



INTERNATIONAL SUSTAINABLE
DEVELOPMENT OBSERVATORY
SECURITY & INTELLIGENCE DEPARTMENT

Manual on False-Flag Operations and Hybrid Warfare in the Context of Asymmetric Warfare

*Concepts, evidence tiers, legal framework, tactical anatomy and
defensive countermeasures*

AUTHORS Álvaro Rodríguez Santiago Sainz	Published by International Sustainable Development Observatory (ISDO), Security and Intelligence Department isdo.ch
Scope Hybrid threats, deniable operations and strategic attribution	Date October 16, 2026 (second edition, corrected October 22, 2026)

Contents

Abstract.....	4
Executive summary.....	5
1. Introduction, scope and method.....	7
1.1 Purpose and audience.....	7
1.2 Evidence tiers, grades and attribution tests.....	8
1.3 Scope and limits.....	9
1.4 Source verification.....	10
2. Concepts: asymmetric warfare, hybrid warfare and the gray zone.....	11
2.1 What selected works on asymmetric conflict explain and what they leave out.....	11
2.2 Schools and critiques of "hybrid warfare".....	13
2.3 Declared doctrines compared.....	15
3. False flags and their neighbors: a working taxonomy.....	17
3.1 What the sources define.....	17
3.2 A working taxonomy.....	18
3.3 Ruses, perfidy and pretext.....	19
3.4 Cyber false attribution and black propaganda.....	20
3.5 The Kydd (2026) conditions.....	21
4. Law and attribution: what the rules reach and what they leave open.....	21
4.1 The instruments and what they cover.....	21
4.2 Deception in armed conflict: where ruses end.....	22
4.3 Force, responsibility and proxies.....	23
4.4 Mapping the elements of a deniable operation to legal regimes.....	24
4.5 Thresholds that remain open.....	26
4.6 Attribution in law and in practice.....	26
4.7 The law-to-practice bridge.....	27
5. Tactical framework: anatomy of a deniable operation.....	30
5.1 Twelve open models and what none of them covers.....	30
5.2 Roles: who fills them.....	32
5.3 Factors and conditions.....	33
5.4 Elements at the doctrinal level.....	34
5.5 Coordination: central approval, plural execution.....	35
5.6 Measuring impact.....	36
5.7 Reaction to pressure and externalities.....	38
5.8 A diagnostic matrix and the attribution posture.....	39
6. Verified practice: documented cases and actors.....	40
6.1 Historical anchors: what a documented false flag looks like.....	40
6.2 Cold War Europe: Italy and Turkey, graded with care.....	42
6.3 Attributed operations by actor.....	43
6.4 Masked identity in cyberspace.....	47
6.5 State operations under denial.....	48
7. Unconfirmed narratives and the misuse of the label.....	50
7.1 Four families, graded claim by claim.....	50
7.2 What the label does in argument.....	53
7.3 A disputed article and the evidence ladder.....	54

7.4 The cost of error in both directions.....	54
8. Detection, assessment and countermeasures.....	55
8.1 Standards for judgment and their evidence base.....	55
8.2 Indicators and open-source validation.....	57
8.3 A two-sided procedure for analysts.....	58
8.4 What the countermeasure record shows.....	61
8.5 Pre-bunking, disclosure and overreach.....	62
8.6 Companies and critical-infrastructure operators.....	63
8.7 Spain.....	65
9. Political economy: why deniable operations are chosen, tolerated and rarely answered in kind.....	66
9.1 What the initiating actor gains.....	66
9.2 What attribution costs the target and its allies.....	67
9.3 The EU listing rounds as a case.....	68
9.4 Who decides.....	70
9.5 What could change the calculus.....	70
10. Policy options: the hammer and the carrot.....	71
10.1 The hammer: legal and financial options.....	72
10.2 The carrot: institutional incentives.....	73
10.3 Enablers, sequence and a caution.....	74
11. Conclusions and research agenda.....	75
11.1 Findings.....	75
11.2 Limits.....	76
11.3 Open questions.....	76
11.4 What future editions should add.....	77
Closing note.....	77
Bibliography.....	80
Annex A. Glossary.....	111
Annex B. Case Register.....	115
B.1 Incidents graded on the three tests.....	116
B.2 Judicial and quasi-judicial register.....	129
Annex C. Methodology and source verification log.....	139
C.1 Verification procedure.....	139
C.2 Verification log.....	143
Annex D. Defensive operating protocols.....	152
D.1 Forensic protocols.....	153
D.2 Verification operations guide.....	163
D.3 Incident response in three parallel flows.....	173
D.4 Record templates.....	181
D.5 Limits and reading route.....	184
Annex E. Spanish summary and bilingual glossary.....	186
E.1 Resumen ejecutivo.....	186
E.2 Glosario inglés-español.....	189

Abstract

This manual examines false-flag operations within the wider field of asymmetric and hybrid warfare, for analysts, defenders, journalists, policymakers and security professionals, with attention to Europe and Spain. It sorts evidence into three tiers (declared doctrine, verified practice, unconfirmed narrative), grades cases and claims, and tests attributions for technical traceability, multi-source convergence and evidence of intent. It proposes a working taxonomy of the false flag and its neighbors, compares legal standards with documented practice, and breaks the deniable operation into roles, conditions, coordination, impact measurement and reaction to pressure. Of 38 contemporary incidents from 2018 to 2026 graded on the three tests, 20 meet none, six meet at least two and one meets all three, counting a test as met only when it is graded Y; if a partly met test also counts, 17 meet all three and four meet none. The grades are one reader's judgment, and a blind second coding by both authors is planned. Of fifteen false-flag and related claims, selected within four families and graded against institutional findings, eight are refuted by a body that examined the specific claim; the set is no sample. Detection standards exist, but the sources read show no outcome evaluation of any institutional countermeasure, so the manual offers a two-sided procedure with explicit evidential thresholds. It explains why deniable operations persist and rates twelve policy options for political viability. The evidence cutoff is October 11, 2026. The second edition adds Annex D, defensive protocols for preserving and testing digital evidence, for the first 24 hours, 72 hours and 30 days after an incident comes to light, and for a response in three parallel flows; they are the authors' proposal and have not been tried on a live incident. The Closing note lists the corrections of October 22, 2026. The manual contains no procedure for staging an operation or evading detection.

Keywords: false flag; hybrid warfare; asymmetric warfare; gray zone; attribution; information manipulation

Executive summary

ISDO's Security and Intelligence Department prepared this manual for readers who must judge claims that an incident was staged, disguised or falsely blamed. It joins three things usually kept apart: the concepts of asymmetric and hybrid warfare, the legal and doctrinal record, and the practice of attribution, with attention to Europe and Spain. Its method is evidence tiering. Declared doctrine is Tier 1, verified practice Tier 2 and unconfirmed narrative Tier 3, and each attribution is tested for technical traceability, multi-source convergence and evidence of intent (Section 1). The manual is analytical and defensive, takes no political side and contains no procedure for staging an operation or evading detection. The evidence cutoff is October 11, 2026.

Table 1. Key figures

Measure	Figure	Basis
Contemporary incidents graded on the three tests	38	Register of 2018 to 2026 at 11 October 2026 (Section 6; Annex B)
Incidents meeting all three tests, counting Y only	1	North Korean remote IT-worker scheme (Section 6)
Incidents meeting at least two tests, counting Y only	6	Each reached a court, a seizure or a multi-state record (Section 6)
Incidents meeting no test, counting Y only	20	Of the 38; 12 meet one test and 5 meet two (Table 50)
Same counts if a partly met test (P) counts as met	17 meet all three; 4 meet none	29 meet at least two; the rule is in Annex C.1.2
Event to first public attribution by a government	At most 2 days to about 10 years	Crimea, Salisbury, Olympic Destroyer, TV5Monde (Section 6.5)
Event to final finding or admission	About 7 weeks to about 7 years 9 months	Crimea, admission 2014; Salisbury, inquiry report 2025 (Section 6.5)
Claims graded against institutional findings	15: 8 Refuted, 7 Alleged	Selected set within four families, no sample (Section 7)
Latest ruling located on enemy colors	1947	Skorzeny acquittal, 9 September 1947 (Section 4)
FIMI incidents counted for 2025	540; 65 percent unattributed	EEAS 4th report, March 2026 (Section 8)
EU hybrid-sanctions total at renewal	80 individuals, 20 entities	Council, 8 October 2026; counts unreconciled (Section 9)
Institutional countermeasures with an outcome evaluation	None found in the sources read	Five measures in Table 24; two more have effect studies (Section 8)

Measure	Figure	Basis
Policy options rated for viability	12: 2 High, 6 Medium, 4 Low	Section 10

Source: Authors' summary of figures sourced in the sections and annexes cited. FIMI is foreign information manipulation and interference. Incident counts rest on one reader's grades, and five incidents carry a grade R (source not retrieved) that could move up to three of them into the group meeting at least two tests (Annex C.1.2).

Eight findings stand out.

- **The strict false flag is rare.** Under the evidence rule of Section 3, which requires Tier 2 evidence of design, the documented cases are historical: Mukden, Gleiwitz and, through secondary sources only, Mainila and the Lavon Affair, whose chain of command is disputed. Olympic Destroyer and the 2020 Proud Boys emails are the contested contemporary candidates. Most of the operations reviewed use proxies, contractors and personas and blame no third party; the graded register runs from 2018 to 2026 (Section 6).
- **Attribution comes in layers, and strong evidence is the exception.** Counting a test as met only when it is graded Y, one of 38 contemporary incidents meets all three tests, six meet at least two and 20 meet none; if a partly met test counts, 17 meet all three (Annex C.1.2). A first public attribution by a government came at most 2 days after the event in Crimea and 8 days after it in Salisbury, but 2 years 8 months after it for Olympic Destroyer and about 10 years for TV5Monde. A final finding or admission took about seven weeks in Crimea and about 7 years 9 months in Salisbury (Section 6.5).
- **The law reaches the elements and misses the design.** Each regime applies above a threshold that a deniable operation is built to stay under, and every institution reviewed treats naming a State as a sovereign political decision (Section 4). EU person-based sanctions bypass the threshold by listing individuals and firms.
- **Access, more than rhetoric, separates settled claims from open ones.** Of fifteen claims selected within four families, a set that is no sample, eight are Refuted, each after an inquiry, court or inspection body examined the specific claim, and seven stay Alleged (Section 7).
- **Defenders have standards and little outcome evidence.** ICD 203 and the UK yardstick require alternative hypotheses and calibrated language, yet the one randomized test of structured techniques on practicing analysts that was located found no gain in accuracy (Section 8). Only two measures, the RT and Sputnik ban and pre-bunking, have quantified effect studies, and both are narrow.

- **Companies are not asked to name attackers.** NIS2 sets clocks of 24 hours, 72 hours and one month, and on the text read no stage requires naming the actor. Spain keeps public attribution to a third State with its National Security Council, and its NIS2 bill is not yet law (Section 8).
- **Deniability pays in delay and doubt.** Naming an actor costs the namer, the EU hybrid listings are the one dated record of response, and no source shows that any measure changed an operator's conduct (Section 9).
- **Evaluation comes first.** Section 10 rates twelve options, two of them High. Every option can be delivered as an output and none has an outcome metric, so the manual puts shared vocabulary and outcome indicators ahead of new instruments.

Analysts will find a two-sided procedure with thresholds in Section 8, the full register in Annex B and the source-verification log in Annex C. Annex D turns the procedure into forensic protocols, a verification guide for the first 24 hours, 72 hours and 30 days, and three parallel response flows; it is the authors' proposal and is untested (Tier 3). Section 11 lists five questions that the evidence cannot yet answer. A Spanish executive summary and a bilingual glossary are in Annex E.

1. Introduction, scope and method

1.1 Purpose and audience

This manual gives analysts, defenders, journalists, policymakers and security professionals a common method for judging claims of false-flag operations and for placing them inside the wider field of asymmetric and hybrid conflict, with attention to Europe and Spain. It is analytical and defensive. It describes how deniable operations are organized, measured and answered, at the level of open doctrine, official records and scholarship, and it gives no procedure for staging a deception or an attack or for evading detection or attribution. Its conclusions come from comparing legal standards that States have accepted with documented practice, and it takes no political side.

One argument links the sections. Deniability is built into the operation. Staged incidents and proxy fronts (Sections 3 and 6) stay under the thresholds of most legal regimes (Section 4), meet defenses whose effect has seldom been measured (Section 8) and leave the cost of naming a culprit with the namer (Section 9). Section 10 sets out options, and Section 11 lists the research still missing. State attributions appear as the

attributions of the government, prosecutor or court that made them, with the accused party's reply where the research found one. The evidence cutoff is October 11, 2026.

1.2 Evidence tiers, grades and attribution tests

Three tiers organize the evidence (Figure 1). Tier 1 is declared doctrine: an official document, strategy or institutional statement that acknowledges a capability or method. It does not prove that any operation occurred. Tier 2 is verified practice: primary sources, technical or forensic evidence, peer-reviewed research, verifiable records and converging independent attributions. Tier 3 is the unconfirmed narrative: inference, anonymous sourcing, partisan media, absent traceability or a conspiratorial reading of a real event. Tier 3 material appears only as a disputed claim and never as a fact. The authors' own readings are labeled *Inference* and carry Tier 3. For books, articles and reports (Tables 3, 4 and 27), the Tier records the kind of source and the reading-limit column records how much of it was read, so a Tier 2 work read at abstract level carries no central finding (Section 1.4).




Figure 1. The three evidence tiers and what each can support.

Source: Authors' evidence scheme, Section 1.

Cases carry labels of three kinds (Table 2). A grade attaches to a stated characterization and not to a whole event, so one event can appear in several rows. A status says whether the operation was executed, planned only or only alleged. An attribution to a named actor is tested three ways. Technical traceability (T) asks whether published

artifacts tie the act to that actor. Multi-source convergence (C) asks whether independent bodies reach the same attribution without resting on one another's unpublished claims. Evidence of intent (I) asks whether tasking, payments, communications or statements show that the actor directed or wanted the act. The tests are applied to a register of 38 contemporary incidents from 2018 to 2026 (Annex B.1). Crimea (2014), TV5Monde (2015) and the DNC intrusion (2016) appear in tables graded by tier and grade only.

Table 2. Labels used in the tables of this manual

Label	Values	What it records
Tier	1, 2, 3	Kind of evidence: declared doctrine, verified practice, unconfirmed narrative
Grade	Documented; Contested; Alleged; Refuted	Confidence in one characterization of a case
Status	Executed; Planned only; Alleged	Whether the operation happened, was only planned, or is only claimed
T/C/I	Y, P, N, U or R for each test	Strength of an attribution on the public record

Source: Authors' scheme. Contested means strongly evidenced but contested; Alleged means alleged or unproven; Refuted means refuted or conspiratorial. Y is met on the public record, P partly met, N not met, U not public, R not retrieved by this review, so the evidence may be public.

The test grades are the authors' judgment from the sources read, and no source publishes scores. Each grade is one reader's reading, and a blind second coding by both authors is planned (Annex C.1.2). In summary counts a test is met only when it is graded Y, so P, N, U and R all count as not met. The manual prints that Y-only count first and, beside it, a sensitivity in which a P counts as met. Grade R marks a source the review did not retrieve, so any count that rests on an R cell says so. Two rules govern the grades. An event is classified as a false flag only when Tier 2 evidence of design exists, and otherwise it is alleged and stays in Tier 3 (Section 3). "Refuted" is used only where an institution or court examined the specific claim, and every other claim stays "Alleged" (Section 7).

1.3 Scope and limits

The manual covers concepts, law, the tactical anatomy of deniable operations, documented practice from 1931 to 2026, the misuse of the label, detection, political economy and policy options, with emphasis on Europe and Spain. Annex D adds defensive protocols for digital and open-source evidence, which are the authors' proposal and have not been tried on a live incident. Annex E gives a Spanish executive summary and a bilingual glossary, and the English text prevails where they differ.

Several things are left out. Gender and age disaggregation is not covered in this edition: the evidence base records the acts of States, organizations and charged individuals, and the sources read were not screened for breakdowns by gender or age. Candidates for a later supplementary table are the EEAS reports on foreign information manipulation, whose treatment of targeted groups was not checked, and victim records in the Salisbury and Crocus City Hall cases. The first table planned covers the exposure and belief evidence of Section 7. No other intersectional dimension is covered, and all of it is left to future research. Kilcullen and Rupert Smith were not reviewed in this edition, and this manual says nothing about their positions. Events after October 11, 2026 lie outside the record, so the status of every trial, listing and investigation may change. State denials were not retrieved for most incidents, and tables say where none was found. Many primary documents, treaties included, were read through extracted excerpts or summarizing tools, so quotations and pinpoints need rechecking against the original text before verbatim use (Annex C).

1.4 Source verification

The research began from a commissioning list of 34 sources, 17 filed under Tier 1 and 17 under Tier 2 by the list's own tier labels, together with seven narrative families for Tier 3. Each source was checked for existence, authors, year, venue and findings, and this manual follows the verified record and never the one-line description on the list. The check audited the list as supplied, so its outcomes describe the list and not the manual's own sources. Of the 34 entries, 19 matched the record as listed, 11 matched with a mismatch such as a wrong year, number, author, venue or kind of document, one matched only a candidate that the entry did not name, and 3 were not found (Annex C). Nine of the 15 entries that could not be used as described are not cited in this manual, and the other six are cited from the verified record. One Tier 2 entry and one narrative family concern the same article in the *American Journal of Economics and Sociology*.

A matching record does not make a strong source. Five regional-journal items show no confirmed Scopus, Web of Science or DOAJ indexing, and they carry no central finding. The disputed article just mentioned is cited for no finding at all (Section 7). Wherever a source was read through one route only, at abstract level or through a secondary account, the text says so beside the datum and the item never carries a central finding. A central finding is a conclusion stated in the Abstract, the Executive summary, a Key finding box, Section 11.1 or Table 1. One such statement leans on abstract-level records, the Key finding of Section 8.4 that two countermeasures have quantified effect studies, and it asserts only that those studies exist and are narrow.

Kydd (2026) and Levin (2025), also read at abstract level, appear in none of these places.

2. Concepts: asymmetric warfare, hybrid warfare and the gray zone

False-flag operations have no doctrine of their own. They appear, when they appear at all, inside two larger vocabularies: the theory of asymmetric conflict and the debate over hybrid warfare and the gray zone. This section sets out what each vocabulary explains, how strong its sources are, and how official doctrines name the same terrain. One point carries forward. Nearly every definition below treats difficulty of attribution as a feature of the terrain, and none treats deliberate misattribution as a defining feature.

2.1 What selected works on asymmetric conflict explain and what they leave out

The classical literature answers a question about outcomes: why weaker actors sometimes win and stronger ones lose small wars. Mack (1975) points to asymmetry of interest and resolve, since the external power has "no direct survival interest at stake" (abstract level only) [1]. Arreguín-Toft (2005), in the Belfer Center's summary, points to the interaction of strategic approaches: opposite approaches favor the weak [2]. Münkler (2003) separates asymmetries of strength from asymmetries of weakness [3], and Kaldor (2013) offers "new wars" as ideal types [4]. Mao's *On Protracted War* (1938) is a statement of intent and no evidence of practice, yet it names a deceptive route to initiative: "exploiting the enemy's misconceptions and unpreparedness" [5]. Qiao and Wang (1999), two serving PLA colonels, list asymmetry among eight principles of "beyond-limits combined war". The translation is a primary source for what the authors advocate and carries no official status [6]. Metz and Johnson (2001) define asymmetry as "acting, organizing, and thinking differently than opponents" (p. 5) and cite Sun Tzu on deception [7].

Table 3. Selected works on asymmetric conflict and what they say about deception

Work	Core claim	Deception or attribution content	Reading limit	Tier
Mao, On Protracted War (1938)	Weaker side advances through three stages	Exploit the enemy's misconceptions to gain initiative	Primary text	1
Mack (1975)	Asymmetry of interest and resolve explains outcomes	None in the abstract	Abstract only	2
Arreguín-Toft (2005)	Opposite strategic approaches favor the weak	None in the summary	Publisher summary only	2
Metz and Johnson (2001)	Asymmetry means acting and thinking differently	Quotes Sun Tzu on deception	Publisher page	2
Münkler (2003)	Asymmetries of strength and of weakness	Deception not treated as a strategy	Article text	2
Kaldor (2013)	"New wars" as ideal types	None in the passages read	Article text	2
Qiao and Wang (1999)	Weaker states offset inferiority by unrestricted means	Asymmetry is one of eight principles	Translation; not doctrine	2

Source: Authors' compilation from [1, 2, 3, 4, 5, 6, 7]. Tier 1 for Mao records a declared statement of intent and no verified practice.

None of these works theorizes false attribution. Deniability enters through other literatures. Cormac and Aldrich (2018) argue that technology, media, special forces and private military contractors have made plausible deniability less credible, so that leaders now embrace "implausible deniability" (abstract only) [8]. Eisenstadt (2021) calls Iran a "quintessential gray zone actor" that uses proxies "to preserve deniability", then concludes that the standoff the proxies provided "was apparently more important to Iran than deniability" [9]. Carson (2018), known here only through the publisher's description, argues that leaders keep wartime involvement secret in order to limit a war [10]. Kydd (2026) is the one formal model of false flags found, and the conditions stated in its abstract do not mention relative strength (Section 3) [11].

Two readings follow. The evidence does not support a clean claim that weak actors use false flags more than strong ones, and no empirical comparison was found. And the best-documented cases of pretext, proposal or misreporting involve actors that held decisive material strength over their targets: Germany in 1939, the Soviet Union in 1939 (documented here only through secondary sources), the United States in 1962 (a proposal) and the United States in 1964 (a misreported incident). *Inference*. Those cases were selected because they are documented, so they cannot carry a general claim about strong states either.

2.2 Schools and critiques of "hybrid warfare"

Hybrid warfare has no agreed definition. Three scholarly lines run side by side. Hoffman (2007) describes fused modes of fighting, "conventional capabilities, irregular tactics and formations, terrorist acts ... and criminal disorder", with Hezbollah in 2006 as the clearest example [12]. Mazarr (2015) frames the gray zone as integrated campaigns that use mostly nonmilitary tools, stay under escalatory thresholds and move gradually toward their objectives [13]. RAND's gray-zone report (Morris et al., 2019) defines an "operational space between peace and war" in which coercion works partly by "blurring the line between military and nonmilitary actions and the attribution for events", and it names lack of attributability a third characteristic [14]. A multinational capability project chose "describing, rather than defining" the phenomenon [15].

The critique is direct. Kofman and Rojansky (2015) call the term "inherently imprecise" and argue that the West wrongly raised Russia's operations in Ukraine "to the level of a coherent or preconceived doctrine" [16]. Stoker and Whiteside (2020) judge that the gray-zone and hybrid-war terms "cause more harm than good" (abstract only) [17]. A NATO Review essay of 7 May 2015, which a University of Glasgow repository record attributes to Damien Van Puyvelde, calls the concept "elusive and catch-all" [18]. On Russia, Renz (2016) argues that the label distorts views of Russian military modernization (abstract only) [19], and Renz and Smith (2016) call the hybrid reading of Gerasimov's 2013 article "selective and hindsight-based" (read through a summarizing tool) [20]. Galeotti (2018), who coined "Gerasimov doctrine", now says of it: "It doesn't exist" [21].

Two later works resist the retreat from the label. Person, Kulalic and Mayle (2024), in *International Affairs* 100(4), accept that many features of the 2014 Crimea operation are absent from the war in Ukraine but argue that hybrid frameworks should stay, because limits on escalation between nuclear-armed rivals make "non-military asymmetric tools of will increasingly define great power competition". They propose placing conflicts on three continuous axes, scale, militarization and centralization, with every conflict showing some degree of "hybridity" (abstract only) [22]. Radin (2017), writing for RAND on the Baltics, sorts Russian hybrid aggression into three forms: nonviolent subversion, covert violent action, and conventional aggression backed by political subversion. He judges the first unlikely to succeed in Estonia and Latvia, the second unlikely to succeed alone, and the third the greatest threat [23]. *Inference*. The dimensional proposal suits this manual better than a category label,

because it lets an analyst describe an operation's profile without first deciding whether the operation is "hybrid".

Table 4. Schools and critiques of "hybrid warfare"

Source	Position	Type and reading limit	Tier
Hoffman (2007)	Hybrid wars fuse conventional, irregular, terrorist and criminal modes	Think-tank monograph	2
Mazarr (2015)	Gray zone: mostly nonmilitary tools, below escalation thresholds	Abstract page	2
RAND RR2942 (2019)	Gray zone blurs military lines and attribution	Report; corroboration only	2
MCDC (2017)	Synchronized use of many instruments; "describing, rather than defining"	Read in full; disclaims official status	2
Kofman and Rojansky (2015)	Term imprecise; no new form of warfare	Think tank; corroboration only	2
Stoker and Whiteside (2020)	Terms "cause more harm than good"	Abstract only	2
Renz (2016); Renz and Smith (2016)	Hybrid doctrine is a post-Crimea Western construct	Abstract only; summarizing-tool reading	2
Galeotti (2018)	The "Gerasimov doctrine" does not exist	Op-ed by the label's originator	2
Person, Kulalic and Mayle (2024)	Keep hybrid frameworks; use scale, militarization, centralization	Abstract only; peer reviewed	2
Radin (2017)	Three forms of Russian hybrid aggression in the Baltics	RAND report; series-level peer review	2

Source: Authors' compilation from [12, 13, 14, 15, 16, 17, 19, 20, 21, 22, 23].

Inference. Three strands travel under the one word. One is a military-operational description of fused modes of fighting (Hoffman). A second describes strategic competition through threshold management and blurred attribution (Mazarr, RAND, and the US phrase "competition below armed conflict"). A third is a governance concept built to organize resilience and legal thresholds (NATO and the EU). Much of the stretching critique is a complaint that the third strand is being justified with evidence from the first.

Key finding

Attribution runs through nearly every definition reviewed. The EU's 2016 working description speaks of "generating ambiguity", its 2018 Joint Communication of campaigns "designed to be difficult to detect or attribute", RAND of "blurring" attribution, and Spain's 2021 strategy of "dificultad de atribuir su autoría". None makes misattribution by design a defining feature. A false flag is therefore a possible instrument inside hybrid and gray-zone repertoires, overlapping with them, and no constitutive part of them.

2.3 Declared doctrines compared

Doctrines record what governments claim to intend, fear or permit. They show nothing about what any actor did. Table 5 keeps to declared positions and states how each text was read.

Table 5. Declared doctrines compared

Actor	Document and date	Declared position	Reading limit	Tier
Russia	Gerasimov article, 2013	Nonmilitary measures supplemented by "military means of a concealed character"; no word "hybrid"	Institutional statement; Military Review translation	1
Russia	Military Doctrine, 2014	Non-military measures using protest potential; indirect and asymmetric methods	Single-route; approval date not shown	1
Russia	National Security Strategy, 2021	Information campaigns and "color revolutions" framed as external attacks on Russia	Unofficial translation; single-route; text errors	1
China	Three Warfares (guidelines dated 2005); Qiao and Wang, 1999	Public-opinion, psychological and legal warfare; "beyond-limits combined war"	No PLA primary text read	1, via secondary
Iran	No official text retrieved	Outside analysts describe proxy use and unacknowledged attacks	Eisenstadt (2021) and WINEP (2020) only	None found
NATO	Warsaw 2016; Brussels 2021; Strategic Concept 2022	Target responds first; attribution is national; hybrid operations may reach armed attack	Paragraph numbers as extracted	1

Actor	Document and date	Declared position	Reading limit	Tier
EU	JOIN(2016) 18; Strategic Compass 2022; Council Conclusions 2026	Flexible working description; Hybrid Toolbox; cost imposition; no legal definition	Full texts read	1
United States	Joint Concept for Integrated Campaigning, 2018; NDS, 2026	"Competition below armed conflict"; 2026 excerpt lacks "gray zone" and "hybrid"	Single-route; excerpt only	1
United Kingdom	National Security Strategy, 2025	Russian "indirect and sub-threshold activity"; new state-threat powers	Extract; may be truncated	1
Spain	ESN 2021; Directiva de Política de Defensa 2020	"Estrategias híbridas" hard to attribute; defense directive adds "zona gris"	BOE text read; directive single-route	1

Source: [6, 9, 24, 25, 26, 27, 28, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39].

Three patterns stand out. The Russian framing inverts the Western one. The word "hybrid" is absent from Gerasimov's article, from the 2014 Military Doctrine and from the English translation of the 2021 strategy (the last two read through a summarizing tool), and in the Russian discourse, as Pynnöniemi (2021) reads it, "hybrid war" is what the West does to Russia [40]. A manual that calls Russian practice "hybrid warfare" is borrowing a Western label. That paper comes from Hybrid CoE, a body whose members include NATO and EU states, which matters when it describes Russian discourse.

NATO's wording hardened. Warsaw (2016) placed primary responsibility for response with the targeted nation. The 2022 Strategic Concept and the 2024 Washington declaration add that hybrid operations could reach the level of an armed attack [31, 41]. The 2025 Hague text never uses "hybrid" and the 2026 Ankara text uses it once, but both are short and spending-centered, and reading a policy retreat into keyword counts in two brief declarations would be an overreach. The EU never adopted a legal definition. Its 2016 working description remains the reference, and the March 2026 Council conclusions turn toward cost imposition [34].

Spain's texts sit on two layers. The 2021 strategy defines "estrategias híbridas" by coordinated, multidimensional action, exploitation of vulnerabilities and difficulty of attribution, and a text search of it found no "falsa bandera", "zona gris" or "negación plausible". The 2020 defense policy directive adds "la zona gris por debajo de otros umbrales de respuesta" and was not searched for the false-flag term [38, 39]. The

directive predates the strategy by about sixteen months, so gray-zone wording was in Spanish defense policy before the strategy left it out. Whether that was a choice is undocumented.

Caution

Declared doctrine is not verified practice. The North Atlantic Council statement of 2 May 2024 and the Council conclusions of March 2026 attribute campaigns to Russia, but they are political declarations resting on national assessments. Whether an individual incident occurred as alleged needs Tier 2 evidence: indictments, court records or forensic reports.

None of the documents in Table 5 defines a false flag. Section 3 therefore builds a working taxonomy from the sources that come closest to doing so.

3. False flags and their neighbors: a working taxonomy

Section 2 showed that hybrid and gray-zone vocabularies treat blurred attribution as a feature of the terrain. The phrase "false flag" needs boundaries of its own. It does at least five different jobs, and most disputes about false flags are disputes about which job is meant.

3.1 What the sources define

The primary institutional sources reviewed contain no operative definition of "false flag". The idea is spread across four better-defined families: covert action with a plausible disclaimer, the law-of-war boundary between ruses and perfidy, attribution in cyberspace, and influence activity that conceals or falsifies its source. The 1948 directive NSC 10/2 supplies the first formula, operations so conducted that the government "can plausibly disclaim any responsibility" [42]. The statutory definition in 50 U.S.C. 3093(e) turns on a role that "will not be apparent or acknowledged publicly" [43]. Neither the EEAS report on foreign information manipulation and interference (FIMI) nor Spain's 2021 strategy uses the term [38, 44].

Dictionaries and scholarship do. Merriam-Webster gives three senses: a hostile action designed to look as if someone else carried it out, a deliberate misrepresentation of motives or identity, and a flag that disguises identity, as on a ship. Its usage examples include "false flag conspiracy theory", a small record of the term's rhetorical drift [45]. The one peer-reviewed definition retrieved sits in the abstract of Kydd (2026): an actor "covertly carries out an attack in hopes that an adversary will take the blame for it"

(abstract only) [11]. The San Remo Manual of 1994 is the only legal text in the set that uses the words, and there they mean a warship's colors [46]. Three senses should stay apart: attribution deception, the law-of-war rule on false colors, and the registry sense of a vessel flying a flag it is not entitled to.

NATO and US doctrine have no entry for the term. The nearest concepts are black psychological operations (products that "appear to emanate from a source other than the true one"), deception ("deliberate measures to mislead targeted decision-makers") and covert operations, in which government sponsorship "is denied if exposed" [47, 48, 49, 50]. The two NATO editions read are withdrawn and the US publication dates from 2006, so they show doctrinal lineage and no current wording. *Inference.* A false flag would be filed under black psychological operations if the product is a message, under deception if it is a battlefield act, and under covert action if a state hides its role. No doctrinal term covers the designed ascription of responsibility to a third party.

3.2 A working taxonomy

Five criteria, each answerable from open sources, let an analyst locate an alleged case before arguing about it.

Table 6. Five criteria for locating an alleged case (authors' proposal, Tier 3)

Criterion	Values	Anchor in the sources
Who acts	State organ; proxy or cutout; non-state group; private contractor	NSC 10/2; Cormac and Aldrich (2018)
Domain	Maritime or kinetic; cyber; information; legal or diplomatic	AP I Art. 39; Tallinn Rules 122 to 127; Rid and Buchanan (2015)
Objective	Pretext for action; blame-shifting; deterrence evasion; discrediting a target	1939 cases; Northwoods; RAND RR2942; Eisenstadt (2021)
Audience	Adversary decision makers; own public; third states and bodies; victim society	Rid (2017) on "an adversary's political decisions"
Attribution mode	Concealment only; denial after exposure; misattribution by design; counter-accusation	Rid and Buchanan (2015); Cormac and Aldrich (2018)

Source: Authors' proposal built on [8, 9, 14, 42, 51, 52, 53].

One test per neighbor keeps the categories apart. Covert action asks whether the sponsor's role is concealed. A pretext operation asks whether an incident, real or staged, is used to justify a response already intended. A false flag asks whether blame is placed on a specific third party by design. Black propaganda asks whether the source of a message is falsified. Perfidy asks whether a combatant feigns protected

status to betray the adversary's confidence. Only the third test needs a designated scapegoat, and only the last belongs to the law of armed conflict. One event can pass several tests, which is why the 1939 radio-station episode counts as both a pretext and a false flag.

Three working definitions follow, all proposals. **Proposal A** (core): a false-flag operation is an act, or set of acts, deliberately designed, in the act itself or in the evidence around it, so that responsibility will be ascribed to a different, specific actor, in order to gain a political, military or legal advantage from that ascription. **Proposal B** (subtype): a false-flag pretext operation is a Proposal A operation whose principal function is to justify a response the sponsor intended anyway; it needs evidence of prior intent, which rarely surfaces before archives open. **Proposal C** (umbrella): attribution deception, any deliberate effort to make an audience believe something false about who is responsible for an act or an information product. C also covers black propaganda and planted forensic artifacts, and it invites the stretching criticism made of "hybrid warfare" in Section 2.

Assessment

On the evidence collected, "false flag" is a distinct analytic category and no legal one. Covert action conceals the sponsor, perfidy feigns protected status, and black propaganda falsifies a message's source; none requires that a particular third party be made to look responsible for an act, and that requirement is what a false flag adds. This manual adopts Proposal A, presents B as a labeled subtype, reserves C for the heading "attribution deception", and applies one evidence rule: an event is classified as a false flag only when Tier 2 evidence of design exists. Otherwise it is "alleged" and carries Tier 3 status.

Applied to the register in Annex B, Proposal A admits Gleiwitz and Mukden (documented), Mainila and the Lavon Affair (through secondary sources only, and with a disputed chain of command in the second) and Olympic Destroyer (contested). The Proud Boys emails of 2020 (Section 5.2) also fit the definition on the allegation of the DOJ, which Iran rejects, and stand outside the register as a second contested contemporary case. It excludes Tonkin (a misreported incident), Northwoods (planned only), the naval colors ruse and covert action without staged blame, and it leaves the 1999 apartment bombings unassessed [11, 54, 55, 56].

3.3 Ruses, perfidy and pretext

Law-of-war rules police the exploitation of legal status and say nothing about authorship. Additional Protocol I, Art. 37, prohibits perfidy, acts that invite an

adversary's confidence in protection "with intent to betray that confidence", and says ruses "are not prohibited"; camouflage, decoys, mock operations and misinformation are the listed examples. Art. 39(2) bars use of the adverse party's flags, insignia or uniforms "while engaging in attacks or in order to shield, favour, protect or impede military operations", and Art. 39(3) leaves the rules on flags at sea untouched [57]. San Remo adds that a warship may not launch an attack under a false flag [46]. State practice on enemy uniforms is divided: the United States has said it does not support the Art. 39 prohibition, and a US military court at Dachau acquitted all the accused in the Skorzeny trial in 1947 [58, 59]. The sources read give the evidence and the verdict, and none gives the court's reasoning.

Pretext and misreporting are different things. The documented standard case of a staged pretext is 1939; Tonkin is a real incident that was misreported, and Northwoods is a plan that nothing shows to have been carried out (Section 6) [54, 55]. Collapsing the three inflates the category.

3.4 Cyber false attribution and black propaganda

In cyberspace the false-flag problem becomes a problem of evidence. Rid and Buchanan (2015) call attribution "an art as well as a science" and note that language artifacts are "relatively easy to fake in a sophisticated 'false flag' operation"; the quotations come from the authors' workshop version, whose pagination differs from the journal [51]. The best-documented planted-artifact case is Olympic Destroyer (February 2018). A vendor found a copied metadata artifact that matched a sample linked to the Lazarus group yet conflicted with the binary's own build characteristics, and it called the pattern an "intricate false flag operation" without naming a group [60]. In October 2020 the United Kingdom said that GRU Unit 74455 "attempted to disguise itself as North Korean and Chinese hackers", and the United States announced an indictment of six officers; the American release contains no false-flag language [56, 61]. Capability borrowing is a neighbor and no twin. A joint NSA and NCSC advisory of 21 October 2019 judged that Turla used tools and infrastructure of an Iranian actor whose operators were "almost certainly not aware of, or complicit with" the use, and the advisory itself does not use the term [62].

Black propaganda is the information-domain cousin. NATO's older psychological-operations doctrine sorted products by attributability into white (acknowledged), gray (source not revealed) and black (apparently from another source), and held that "black" campaigns are generally less successful than white ones [47]. That is a stated doctrinal view from a withdrawn edition, with no test attached.

3.5 The Kydd (2026) conditions

Kydd (2026) offers the only formal model of false flags found. This manual read its abstract only, because the article is closed access. The abstract opens with the Soviet shelling of Mainila in 1939, announces a game-theoretic model and states the conditions under which a false flag could occur in equilibrium: the actor does not find it too costly, the adversary wants to carry out that kind of attack, and a third party is willing to punish the adversary because it is more likely to be guilty [11]. It names the 1999 Russian apartment bombings only as "a potential case", and none of its stated conditions mentions relative strength. *Inference.* If the full article bears out the abstract, the structural prerequisites are an audience that can be steered and a scapegoat that fits the pattern; weakness as such drops out. That reading needs checking against the article's propositions before anyone relies on it.

One empirical companion exists. Levin (2025) reports survey experiments in which third-party publics became more likely to approve of war-seeking states after low-credibility false-flag incidents and less likely to back sanctions, while many domestic publics, the American one included, found such incidents ineffective or counterproductive. The abstract ties that result to low trust in one's own government; sample sizes and countries were not seen, and the abstract was read on the author's own page [63]. The payoff of a false flag, on this evidence, depends on audience trust.

4. Law and attribution: what the rules reach and what they leave open

Sections 2 and 3 treated the false flag as an analytic category: a deniable act designed so that blame lands elsewhere. This section asks what the law does with it. Treaty and customary rules regulate the conduct around a false flag (deception in war, force, intervention, State responsibility), and each applies above a threshold that a deniable operation is built to stay under. The act that would trigger most of those rules, public attribution, is treated by every institution reviewed as a sovereign political decision with no duty to disclose evidence. Legal texts and declared positions are Tier 1, rulings and official registers are Tier 2, and the authors' own readings are Tier 3 (Section 1).

4.1 The instruments and what they cover

Only one text in Table 7 contains the words "false flag": paragraph 110 of the San Remo Manual (1994), a non-binding restatement for naval warfare [46]. Every other

instrument regulates something next to the deception, such as insignia, force or the conduct of a State's organs.

Table 7. Legal instruments that bear on false flags and hybrid operations

Instrument	Provisions	Relevance to deniable operations	Tier
Hague IV Regulations (1907)	Arts 22, 23(b), 23(f), 24	Bans treachery and improper use of enemy insignia; permits ruses	1
Additional Protocol I (1977)	Arts 37-39	Bans perfidy; ruses lawful; neutral and enemy colors	1
ICRC custom study (2005); San Remo Manual (1994)	Rules 57-65; paras 110-111	Non-binding restatements; San Remo alone names "false flag"	1
Tallinn Manual 2.0 (2017)	Rules 66, 122-127	Non-binding expert text on intervention, cyber perfidy, ruses	1
UN Charter (1945)	Arts 2(4), 2(7), 39, 51	Force, domestic jurisdiction, aggression, self-defense	1
ILC Articles on State Responsibility (2001)	Arts 4, 8, 11, 25, 49-54	Attribution of conduct, necessity, countermeasures	1
Rome Statute (1998); Kampala amendments (2010)	Arts 8(2)(b)(vii), (xi); 8 bis	Insignia misuse, treachery, aggression; text not re-read	1
UN GGE (2015, 2021); EU Toolbox (2017); Council Declaration (2024)	A/70/174 paras 13, 28; Toolbox paras 4-5	Attribution is sovereign and political; no proxies	1
Decisions (CFSP) 2019/797 and 2024/2643	Arts 1, 4; Arts 1(1)(a), 2	Person-based sanctions: cyber-attacks, FIMI, sabotage, migration	1, 2
NATO texts 2014-2026	Wales 72; Warsaw 70, 72; Brussels 31-32	Hybrid and cumulative-cyber thresholds; attribution a national prerogative	1

Source: Authors' compilation from [29, 30, 46, 53, 57, 58, 64, 65, 66, 67, 68, 69, 70, 71]. Rome Statute provisions are cited as identified, with the primary text not re-read [72].

Tallinn, San Remo and the ICRC study are expert restatements and no treaty. Most treaty texts were read through excerpts, so pinpoints need checking against official PDFs before verbatim quotation (Annex C).

4.2 Deception in armed conflict: where ruses end

Section 3 set out the ruse and perfidy line. Three timing rules decide when the use of enemy colors crosses it, and they differ. The Hague Regulations bar the "improper use" of the enemy's flag, insignia and uniform and never define "improper" [64]. Additional Protocol I narrows the ban to the moments when a party is "engaging in attacks" or acting to shield, favor, protect or impede military operations [57]. At sea, San Remo

allows a false flag until the attack begins [46]. A statement that false flags are prohibited in war has to name which of the three it relies on.

Cyberspace adds a split among experts. The Tallinn experts divided on whether the confidence that perfidy exploits must be human, and the commentary to Rule 122 says the rule does not extend to perfidious acts that result in damage or destruction of property [73, 74]. Price's case for a human-confidence requirement appears in an online law-review piece not confirmed as peer reviewed [74].

Adjudicated practice is thin. After the Skorzeny acquittal of 9 September 1947 (General Military Government Court of the U.S. Zone of Germany, Dachau, no docket number shown, final), the review found no ruling on enemy uniforms, flags or perfidy before the ICC, the ad hoc tribunals or any national court [59]. That is 79 years, or 49 since Additional Protocol I. The search was not exhaustive, so the finding is an absence in the sources reached.

Inference. These rules police the exploitation of legal status and say nothing about authorship. An operation that plants artifacts pointing at a third State feigns no protected status and invites no confidence in IHL protection. The perfidy and ruse rules do not reach it, and without an armed conflict they do not apply at all (Tier 3).

4.3 Force, responsibility and proxies

Charter Art 2(4) bars the threat or use of force, and Art 51 permits self-defense only "if an armed attack occurs" [65]. The Council of the EU's November 2024 Declaration places cyber operations comparable to a kinetic use of force under Art 2(4) and allows that the combined effects of several operations could, taken together, be comparable [69].

Proxies are where attribution becomes a legal test. ILC Art 8 attributes the conduct of a person or group "in fact acting on the instructions of, or under the direction or control of" a State, and the commentary confines this to cases where the State "directed or controlled the specific operation" [66, 75]. In *Nicaragua v. United States* (ICJ, merits, 27 June 1986, para 115) the Court required proof of "effective control of the military or paramilitary operations" in which the violations occurred, and still found the United States in breach of the duties not to intervene and not to use force through its own conduct [76, 77]. In *Bosnia and Herzegovina v. Serbia and Montenegro* (ICJ, 26 February 2007) the evidence did not prove that the acts of the Bosnian Serb army could be attributed to Serbia, although Serbia breached its duties to prevent and to punish genocide; the remedy was declaratory [78]. The Court also rejected the Tadić

"overall control" test as wider than the ILC's three standards, a passage this review knows only through Cassese (2007) [79].

The ECtHR Grand Chamber took another route on 9 July 2025. In *Ukraine and the Netherlands v. Russia* (nos. 8019/16, 43800/14, 28525/20, 11055/22) it held that Russia had jurisdiction over separatist-held areas from 11 May 2014 until at least 26 January 2022 and answered for the acts and omissions of its forces and of the separatist groups. On MH17, "the only reasonable conclusion" absent Russian information was that a Buk crew member or a "DPR" member fired the missile [80]. That is attribution through effective control of territory. ILC Art 8 asks about control of a specific operation, so the two tests cannot be swapped.

Inference. The ICJ record sets a high bar for pinning a proxy's act on a State and a lower one for the State's own support to armed groups. In *Armed Activities on the Territory of the Congo* (DRC v. Uganda) the Court found in 2005 that Uganda's "military, logistic, economic and financial support to irregular forces" breached the non-intervention duty [81]. Arrangements built to stay below "specific operation" control therefore tend to engage the sponsor's own duties and leave the proxy's acts unattributed (Tier 3).

Acting on a wrong attribution carries its own risk. Countermeasures lie only "against a State which is responsible for an internationally wrongful act" (ILC Art 49(1)), and a State that "has contributed to the situation of necessity" cannot plead it (Art 25) [66]. A State acting on an unfounded assessment "may incur in international responsibility for its own wrongful conduct" [82].

Individual criminal responsibility for aggression depends on the forum. The ICC has had jurisdiction under the Kampala amendments since 2018, after 30 ratifications, and its Assembly agreed in July 2025 to meet again in 2029 on a proposed amendment [83]. A tracking site counts "49 States Parties" while the table on the same page lists 48, so the count is given with that caveat [84]. For Ukraine, the Council of Europe adopted an Enlarged Partial Agreement on 15 May 2026 for a Special Tribunal; it reports 38 States and the EU as having notified intent to join, whereas the ICDS counts 34 member States plus Australia, Costa Rica and the EU. The agreement had not entered into force and no indictment exists [85, 86].

4.4 Mapping the elements of a deniable operation to legal regimes

No single regime covers a deniable operation. IHL needs an armed conflict, the force rules need effects comparable to kinetic force, non-intervention needs coercion of a free choice, and State responsibility needs instruction, direction or control. Table 8

applies those thresholds to seven elements that recur in the sources on hybrid and false-flag operations.

Table 8. Elements of a deniable operation, legal regime and EU sanctions hook

Element	Regime and threshold	Attribution route	EU hybrid-sanctions hook (Decision 2024/2643)	Tier
Ruses versus perfidy	IHL, armed conflict only; perfidy needs feigned protected status and intent to betray	Organs (ILC Art 4); individual liability, Rome Statute Art 8(2)(b)(vii), (xi)	None	1
Psychological operations and propaganda	Misinformation is a lawful ruse; ICCPR Art 20(1) requires a legal ban on war propaganda	Units and organs (Art 4); due diligence not engaged	Art 1(1)(a)(iv): coordinated information manipulation and interference	1
Information manipulation (FIMI)	Non-intervention applies only if coercive; persuasion alone falls short on both expert views	Art 8 for contractors and outlets; due diligence disputed	Art 1(1)(a)(i) elections and (iv) information manipulation; listings 2024-2026	1, 2
Covert action through proxies	Art 8 needs control of the specific operation; the sponsor's own support engages other rules	GGE 2015: no proxies; no knowing use of territory	Art 1(1)(a): responsible for, supporting or benefiting; GRU Unit 29155 listed 16 December 2024	1, 2
Cyber false attribution	No rule targets deception about authorship; the act itself is tested under sovereignty, intervention, force	Territorial origin alone insufficient; a wrong response risks the responder's own responsibility	Decision 2019/797 Art 4: persons responsible for attacks; no State finding needed	2, 3
Sabotage of infrastructure	Force if scale and effects match kinetic ones; otherwise sovereignty and criminal law	Eagle S: individual crew; Cook Islands flag; no State link in the source	Art 1(1)(a)(v): economic activities or services of public interest	1, 2
Coercion below the force threshold	Non-intervention requires coercion of a free choice; responses are retorsion or countermeasures	Arts 4 and 8; due diligence not engaged	Art 1(1)(a)(vi): instrumentalization of migrants	1

Source: Authors' mapping from [57, 58, 66, 67, 70, 71, 87, 88, 89, 90, 91, 92, 93]. Rome Statute provisions cited as identified, text not re-read [72].

Inference. The elements most often called false flags in public debate (planted artifacts, deniable proxies, disinformation) sit under the threshold of every hard-law regime except sanctions. EU person-based sanctions bypass the State-responsibility threshold because they list individuals and entities for what they did, and the listing

grounds cite information manipulation and cyber activity without any finding of intervention (Tier 3).

Eagle S shows the sabotage row in practice. The Helsinki Court of Appeal, in a ruling reported on 27 August 2026, held that Finland has jurisdiction and that the events of 25 December 2024 were "not a maritime accident" under UNCLOS, after the District Court had dismissed the case in October 2025. The deadline to appeal is 26 October 2026, so the ruling can still be challenged, and the cable's owners seek about EUR 105 million. The source records individual crew and the Cook Islands as flag State, with no link to any State, and gives no docket number [91].

4.5 Thresholds that remain open

The unsettled questions cluster where deniability works. On sovereignty, the Tallinn 2.0 experts agreed only that a rule exists and is violated by physical effects on another State's territory or by interference with inherently governmental functions; the United Kingdom denies a standalone prohibition (2021, para 10), while the Council's 2024 Declaration treats a violation as wrongful even without intervention or force [69, 88, 94]. Ten parties, among them the Netherlands, France, Germany and the African Union, treat due diligence as binding or customary (2019-2024); five, among them the United Kingdom and the United States, have not accepted that (2020-2021) [95]. On intervention, the United States said in 2020 that "there is no international consensus among States on the precise scope or reach of the non-intervention principle"; no State statement located says that disinformation alone is intervention, and the positions found concern election manipulation [87].

Reform channels are all soft law. Tallinn 3.0 was launched in 2021 as "a five-year venture", and no sign of publication turned up by 11 October 2026 [94, 96]. The Oxford Statement of 2 June 2021 (111 signatories listed) applies sovereignty, non-intervention and IHL to information operations [97]. The UN Global Mechanism, whose first substantive session ran from 20 to 24 July 2026, is meant to produce outputs that are "consensus-based, non-binding, and persuasive" [98, 99]. The one peer-reviewed study located on a non-kinetic psychological tool, sonic booms from military aircraft, finds that IHL "does not explicitly regulate" their use as a tool of terror, and it is silent on false flags, so it is cited only as an example of the regulatory gap [100] (Annex C.2.6).

4.6 Attribution in law and in practice

Every institution reviewed treats attribution as a sovereign political act. The EU's 2017 draft conclusions call it "a sovereign political decision based on all-source intelligence",

and the 2024 Declaration lets each State decide whether to keep an assessment confidential [68, 69]. The United Kingdom sees "no general legal obligation" to disclose underlying information (2021, para 15), Germany requires "a sufficient level of confidence", and NATO's Brussels communiqué calls attribution "a sovereign national prerogative" [30, 88, 89]. The GGE adds a caution: origin in a State's territory "may be insufficient in itself to attribute the activity to that State", and accusations "should be substantiated" [67].

The evidentiary ladders differ. The UK's probability yardstick (24 March 2025) maps "highly likely" to about 80 to 90 percent and "almost certain" to about 95 percent up to just under 100, with a separate judgment of analytical confidence [101]. The ECtHR used "beyond reasonable doubt" for some findings and "the only reasonable conclusion" for MH17 [80]. The ICJ's "fully conclusive" standard for charges of exceptional gravity (Bosnia, para 209) is cited as identified in the research and not re-read. Olympic Destroyer shows why the ladder matters (Section 3): the UK's "almost certainly (95%+)" attribution of 19 October 2020 came two years and eight months after the incident of 9 February 2018 [56, 102].

What the practice produced is a sanctions record. The EU cyber regime stood at 27 individuals and 11 entities at the Council's review of 13 July 2026 [103]. The hybrid regime was renewed on 8 October 2026, to 9 October 2027, with 80 individuals and 20 entities, although a listing of 24 September 2026 is recorded as entry 81; the review could not reconcile the two figures [104, 105]. Section 9 returns to the listing rounds.

Private organizations face a different logic. NIS2 reporting clocks (24 hours, 72 hours, one month) and the sanctions-breach offenses attach to the incident and the counterparty and never to the attacker's identity; the early warning only has to say whether unlawful or malicious acts are suspected [106]. Section 8 returns to that exposure.

4.7 The law-to-practice bridge

Table 9 pairs each obligation with the latest measurable datum the review found, and one rule fixes both clocks. The first runs from the date the obligation was set to the date of the datum. The second runs from the datum to the evidence cutoff of 11 October 2026 and gives the age of the datum. Where no datum was located, the table says so and the first clock runs to the cutoff. Two rows start at an incident and not at a norm, Olympic Destroyer and Eagle S. The data are of five kinds: court rulings, treaty steps, sanctions listings, infringement steps and state or alliance positions. Several are

enforcement events and not measures of compliance. The first interval grows with the age of the norm whatever its record, so it ranks no institution. The age of the datum does not depend on the age of the norm, and Figure 2 charts it.

Table 9. Obligation, latest measurable datum, elapsed time and age of the datum

Obligation (date set)	Latest measurable datum (date)	Obligation to datum	Datum age at cutoff	Tier
Hague 1907 Art 23(f) (18 Oct 1907); AP I Arts 37, 39 (8 Jun 1977)	Skorzeny acquittal, 9 Sep 1947; no later ruling located	39y 10m; AP I: no datum, clock open 49y 4m	79y 1m	2
Charter Art 2(4) (signed 26 Jun 1945); Res 3314 (14 Dec 1974)	Special Tribunal agreement 25 Jun 2025; Enlarged Partial Agreement 15 May 2026; no indictment	80y 10m; 51y 5m	4m 26d	2
Rome Statute Art 8 bis (11 Jun 2010)	48 or 49 States Parties at 7 Sep 2025; jurisdiction active since 2018	15y 2m, to the 2025 count	1y 1m	2
ILC Art 8; Nicaragua para 115 (27 Jun 1986)	ECtHR Grand Chamber, 9 Jul 2025: territorial effective control instead of Art 8	39y 0m	1y 3m	2
Non-intervention (1970, 1986), clock from Decision 2024/2643 (8 Oct 2024)	EU hybrid listing 24 Sep 2026; renewal 8 Oct 2026, 80 individuals, 20 entities	2y 0m	3d	2
GGE 2015 norms (22 Jul 2015); EU cyber regime (17 May 2019)	Listings 13 Jul 2026: 27 individuals and 11 entities; renewed to 18 May 2027	10y 11m; 7y 1m	2m 28d	2
Attribution of planted artifacts (GGE 2015, para 13(b)); clock from the incident of 9 Feb 2018	Olympic Destroyer: UK "almost certainly (95%+)" and US indictment, 19 Oct 2020	2y 8m, incident to attribution	5y 11m	2
NATO hybrid and cumulative-cyber thresholds (9 Jul 2016; 14 Jun 2021)	Article 4 request 10 Sep 2025; Article 5 only in 2001; Ankara declaration of 8 Jul 2026 silent	9y 11m to Ankara; 4y 2m to request	3m 3d	2
Directive (EU) 2024/1226 (transposition due 20 May 2025)	Notices to 18 Member States 24 Jul 2025; Ireland reasoned opinion reported 6 Oct 2026	2m 4d; 1y 4m 16d	5d	2

Obligation (date set)	Latest measurable datum (date)	Obligation to datum	Datum age at cutoff	Tier
NIS2 (transposition due 17 Oct 2024)	Reasoned opinions to 19 States 7 May 2025; Court referral of four, 8 Jul 2026	6m 20d; 1y 8m 21d	3m 3d	2
Cable sabotage, UNCLOS Art 113 as identified (1982); clock from the incident of 25 Dec 2024	Eagle S: Helsinki Court of Appeal affirms jurisdiction, 27 Aug 2026	1y 8m, incident to ruling	1m 14d	2

Source: Authors' compilation from [56, 59, 61, 64, 70, 80, 83, 84, 85, 91, 103, 104, 105, 107, 108, 109, 110, 111, 112, 113, 434, 435, 436]. Dates calculated to the day by the authors under one rule. The first interval runs from the date the obligation was set, or from the incident where the first column says so, to the date of the datum. The age of the datum runs from the datum to the cutoff of 11 October 2026, and where a row holds two data it is the age of the later one. Additional Protocol I has no datum, so its clock runs to the cutoff, and "none located" records the limit of the search. Figure 2 shows the age of each datum in years, as days divided by 365.25. The exact dates are in the open data (bridge_intervals.csv). The days on which Hague IV [64], the Charter [434], Resolution 3314 [435], Protocol I [436] and Decision (CFSP) 2019/797 [70] were adopted or signed are those printed on the pages cited.




Figure 2. Age of the latest measurable datum at 11 October 2026, in years (axis cut at 7 years).

Source: [56, 59, 80, 84, 85, 103, 104, 109, 111, 113].

Key finding

The age of the latest datum, unlike the interval from obligation to datum, does not depend on the age of the norm. Where an institution keeps a register, the latest datum is days to three months old at the cutoff: 3 days for the renewal of the hybrid-threats regime, 5 days for the Ireland reasoned opinion on the sanctions-crime directive, 2 months 28 days for the cyber listings and 3 months 3 days for the NIS2 Court referral. In the other rows the latest datum is between 1 month 14 days and 5 years 11 months old, except for perfidy and enemy colors, where the latest ruling located dates from 1947 and is 79 years 1 month old. A fresh datum shows that an institution recorded something recently and says nothing about compliance.

Inference. Compliance data exist where an institution keeps a register. Where only courts can enforce, the record is sparse, and "no datum located" is the honest entry; it must never be read as compliance. The deniability logic that Sections 3 and 6 trace through staged incidents and proxy fronts therefore recurs in law as threshold design. Section 8 turns to detection, and Section 10 to the options that could change the cost of ambiguity (Tier 3).

5. Tactical framework: anatomy of a deniable operation

Section 4 showed that the law reaches a deniable operation element by element. This section decomposes the operation itself: the models that describe it, the roles that carry it, the conditions that shape it, how it is coordinated, how its impact is measured and how it reacts to pressure. Almost every source is a government or a prosecutor describing an adversary, so each structure below is an attribution by the cited authority and each indictment is an allegation; replies from the accused are given where the research found them. Every case is also one that was detected, and nearly every one was publicly attributed by some authority, so undetected operations cannot be observed here. Procedures are omitted by design.

5.1 Twelve open models and what none of them covers

Twelve open frameworks sit at four levels: system and campaign models, technique-tagging taxonomies (most ported from cyber defense), triage lenses and own-force doctrine. Stacking them requires naming the level at which each applies.

Table 10. Twelve open frameworks by level, what they break down and their stated limits

Framework	Level	Breaks down	Stated limit	Tier
JRC and Hybrid CoE model (2021)	System	Thirteen domains; priming, destabilization, coercion	Gives "no estimation as to if or when"	2
MCDC Countering Hybrid Warfare (2019)	System	Goals, response thresholds, detect-deter-respond cycle	"Generic"; "further work is required"	2
EEAS FIMI reporting (2026)	Technique	Incidents, observables, techniques, manipulation sets	"Selective, time-bound sample"	2
DISARM	Technique	Adversary stages from planning to assessing effectiveness	Defender side "less mature than Red"	2
MITRE ATT&CK (2020)	Technique	Tactics, techniques, groups, mitigations	Attribution "solely based on" techniques not advised	2
Diamond Model (2013)	Technique	Adversary, capability, infrastructure, victim	Adversary field "likely to be empty"	2
Intrusion kill chain	Technique	Seven intrusion phases; defender's matrix	Rests on one illustrative case	2
François ABC (2019)	Triage	Actors, behaviors, content	Platform information asymmetry "immense"	2
ABCDE (NATO 2024; UNHCR)	Triage	Adds degree and effect	Origin of D and E not verified	1
Wardle and Derakhshan (2017)	Triage	Agent, message, interpreter	Exact influence cannot be measured	2
JP 3-13.4 (2006)	Own force	Objective, target, story, channels, feedback	Successor not publicly released	1
AJP-10.1 (2023); AJP-10	Own force	Information operations as a staff function	Cognitive effect "not directly observable"	1

Source: Authors' compilation from [49, 114, 115, 116, 117, 118, 119, 120, 121, 122, 123, 124, 125, 126, 127]. All items were read through a summarizing tool, so short quotations need rechecking.

Three gaps stand out. Only DISARM lists assessing effectiveness as a stage of the adversary's work. None of the twelve has a slot for deliberate misattribution; the nearest items are MCDC's ambiguity, the Diamond Model's split between adversary operator and customer, and the JRC model's threshold manipulation. And each states or implies that it has not been validated against outcomes. The technique-level models came from intrusion analysis, and MITRE itself warns against attributing activity "solely based on the ATT&CK techniques used" [119]. Technique tagging of

influence activity, as in the EEAS reports, is therefore a descriptive device and proves no authorship.

5.2 Roles: who fills them

Six roles recur: principals, executors, proxies and personas, amplifiers, enablers and target-side actors (Table 11 and Figure 3).

Table 11. Six recurring roles and documented holders (attributions by the cited authorities)

Role	Documented holders	Source type
Principals	Russian Presidential Administration and GRU; Chinese MPS and MSS; Iranian Guardian Council and IRGC	Indictments, Treasury notices
Executors	SDA, Structura, ANO Dialog; Internet Research Agency; i-Soon; Emennet Pasargad	Affidavit, indictments
Proxies and personas	Proud Boys persona; DCLeaks and Guccifer 2.0 personas; contractors and criminals	DOJ (2021); Mueller report
Amplifiers	Recruited or unwitting third parties; subcontracted amplification networks	Mueller report; EEAS 4th report
Enablers	US-issued payment cards, an exchanger, registrars, registries, hosting	Doppelganger affidavit (2024)
Target-side actors	Platforms, registries, one newspaper, journalists given false bylines	Doppelganger affidavit (2024)

Source: Authors' compilation from [116, 128, 129, 130, 131, 132, 133, 134, 135, 136, 137].

The executor layer is commercial in the best-documented cases. The Doppelganger affidavit records a translation-invoice trail at SDA (27 invoices, para 62), the DOJ describes a price list at i-Soon, and the EEAS calls "FIMI-as-a-service" a "multi-layered business ecosystem in which activities are planned, financed, and subcontracted" [116, 129, 134].

One device in the record works like a false flag in the narrow sense. According to the DOJ, two contractors of Eeleyanet Gostar, now Emennet Pasargad, sent threatening emails to voters in October 2020 "claiming to be a 'group of Proud Boys volunteers'" and produced a video with the Proud Boys logo that purported to show the hacking of state voter sites; a footnote says the intrusions shown were simulated [136]. Iran's Foreign Ministry called the accusations "baseless" on 19 November 2021 [138]. Most other cases in the research are deniability by distance. Five cyber cases borrow or invent an identity (Table 19), and of those only Olympic Destroyer meets Proposal A, with TV5Monde undecided (Section 6.4; Tier 3).

Principal-side replies are scarce. Prigozhin, through Concord's press service on 7 November 2022, said "we are interfering and we will continue to interfere", which Reuters called the first such admission from a figure Washington had formally implicated [139]. No Russian reply to the 2024 Doppelganger action or to the EU listings was retrieved.




Figure 3. Roles in a documented deniable operation and the flow of tasking, output and feedback.

Source: [116, 128, 129, 137].

5.3 Factors and conditions

The evidence sorts conditions into what the operator can afford, what the target brings, what the environment supplies and what the observer can see. Table 12 labels each as shown (documented in a primary record or measured), argued (a peer-reviewed argument not tested on these operations) or hypothesis (stated without a test).

Table 12. Conditions that shape whether and how far operations go, with evidence status

Condition	Evidence	Status	Limit
Attribution cost falls as target value rises	Lindsay (2015), on cyberattack	Argued	No test on these operations
Deniability is eroding	Cormac and Aldrich (2018)	Argued	Abstract level only

Condition	Evidence	Status	Limit
A false flag needs a plausible culprit and a willing punisher	Kydd (2026)	Argued	Abstract level only; silent on relative strength
Standoff may outweigh deniability	Eisenstadt (2021), on Iran	Argued	One analyst's reading of one actor
Country resilience profiles	Humprecht, Esser and Van Aelst (2020)	Shown at country level	No test against campaign outcomes
Overt interventions beat covert ones	Levin (2016), elections 1946-2000	Shown, narrow	Data end in 2000
Coupling to dated events	Viasat attack "one hour before" the invasion; Olympic Destroyer timed to the opening ceremony	Shown for timing	Intent to maximize effect is an inference
Generative AI	Goldstein et al. (2024); EEAS: 147 AI-related incidents in 2025, 41 in 2024	Shown for use	Survey experiment; organic engagement "remains low"
"Firehose" repetition model	Paul and Matthews (2016)	Hypothesis	Conceptual review; no real-world effect measured

Source: Authors' compilation from [8, 9, 11, 116, 131, 140, 141, 142, 143, 144, 145, 146].

Conditions that make an operation feasible and cheap (attribution cost, timing hooks, tooling) are better documented than those that make it effective; the first set comes from indictments and sanctions notices, the second is mostly theory. *Inference*. Read together, Levin, Lindsay and Cormac and Aldrich suggest a trade between deniability and effect: covertness may cost persuasive power, and denials are believed less anyway. That would fit operators who kept going after being named (see the discussion of adaptation below), but it is a reading and no finding, and Levin's data stop in 2000 (Tier 3).

5.4 Elements at the doctrinal level

Doctrine names a few elements that a defender can look for. JP 3-13.4 lists an objective, a target, the perceptions required of it, a story, the means that carry it, an event schedule, feedback channels, measures of effectiveness and a termination concept, and it insists that the target must act and not only believe [49]. NATO's AJP-10.1 treats information operations as "a staff function to analyze, plan, assess and integrate information activities" [126]. These are frameworks for a State's own planning, and the procedures behind them are left out; what carries over is the analytic questions of Table 15. JP 3-13.4 also prohibits perfidy and bars deception aimed

at the American public, Congress or news media, which puts the line between a lawful ruse and an unlawful act where Section 4 drew it [49].

5.5 Coordination: central approval, plural execution

In every documented case approval runs up a hierarchy while execution is plural: several firms, units or contractors working for one principal, and sometimes two principals working one target in one period (Table 13).

Table 13. Coordination patterns and strains on record (allegations and attributions by the cited authorities)

Case	Approval	Execution	Strain or failure on record
Russia, Presidential Administration (2024)	20+ meetings; leadership reactions recorded	SDA, Structura, ANO Dialog	"Impossible task" note; five countries cut down
Russia, GRU line (same 2024 cycle)	GRU direction and financing, per Treasury	Center for Geopolitical Expertise	No deconfliction between the two lines found
China, 912 group and contractors	MPS headquarters sets topics	MPS officers; i-Soon sells to 43 or more bureaus	Speculative hacking sold on; no deconfliction found
Iran, contractors behind personas	State direction attributed months after claims	Contractors using borrowed identities	Lag between claim and attribution

Source: Authors' compilation from [129, 130, 131, 134, 135, 136, 147].

The widest gap between sources is Russia's. A sworn affidavit of 4 September 2024 records more than 20 meetings with the Presidential Administration (para 49) and the leadership's reactions to work product, such as "well received" and "need to work" (para 51) [129]. A 2017 scholarly account describes an "adhocracy" with no master plan, in which actors compete to show they can serve broad objectives [147]. Prosecutors describe chains of command because charging documents need them, and scholars describe behavior over time, which may explain part of the gap. The affidavit is an allegation resting on the contractor's own documents as quoted by an FBI agent, and Galeotti's paper rests on off-the-record interviews and predates the 2022 operations.

Two principals worked the 2024 US election cycle. Treasury tied SDA and Structura to the Presidential Administration on 20 March 2024 and, on 31 December 2024, tied the Center for Geopolitical Expertise to the GRU, acting "at the direction of, and with financial support from, the GRU" [130, 131]. Two lines in one cycle are documented. Deconfliction between them is not.

China's record is the most formal. The complaint of 17 April 2023 says the topics of the "912 Special Project Working Group", 34 MPS officers, "are directed by MPS headquarters" [135]. The i-Soon case (5 March 2025) adds a contractor layer: sales to at least 43 MSS or MPS bureaus in at least 31 provinces and municipalities at roughly USD 10,000 to 75,000 per inbox, and a DOJ statement that private firms "obscured the PRC government's involvement" [134, 148]. A Chinese embassy spokesperson called the allegations a "smear" [149]. Iran's pattern is persona first, with state direction attributed months later (Section 6).

Shared infrastructure blurs principals further. The EEAS reports that "the same amplification networks can be mobilised by different actors if they are contracted by different customers" and that Russia-China cross-posting happened "on rare occasions" [116]. The failures on record fall into four groups: fatigue at the executing layer (participants who had "initially talked him into five countries. Now he says no"), rivalry among principals (the FSB reportedly persuaded the leadership to release material the GRU had gathered), blowback (clumsy hacking and the Le Pen loan "drove Emmanuel Macron to a more hostile position") and the loss of a principal [129, 147]. On the last, The Bell wrote on 4 July 2023 that Prigozhin's Patriot Media Group, which includes the IRA, was closed after the June 2023 mutiny, on largely anonymous sources (Tier 3) [150]. No source describes a mechanism that resolves conflicts between principals. *Inference*. "Centralized or decentralized" is a poor binary here; the sources support a split between who approves (a center), who executes (plural, often commercial), who amplifies (a market) and who is the principal when two lines overlap (Tier 3).

5.6 Measuring impact

Three levels have to be kept apart: outputs (volume, reach, engagement), outcomes (changed belief or behavior) and strategic effects (a policy or election result). Table 14 shows what the record measures at each.

Table 14. Three levels of impact and what the record measures at each

Level	Operators' indicators on record	Defenders' measures	Evidence strength
Outputs	Quotas, subscriber goals, dashboards, leadership approval	Breakout Scale spread; EEAS authenticity of engagement	Documented, mostly by prosecutors and the EEAS
Outcomes	None found	Survey studies: null or small effects in narrow samples	Few and narrow; "no effect" and "alone" both unsupported
Strategic effects	None found	Levin (2016) election data to 2000; Dylan and Maguire (2022)	Sparse; largely inference

Source: Authors' compilation from [116, 129, 135, 137, 142, 151, 152, 153, 154].

Operators' own indicators, where documents show them, are outputs plus the principal's approval. SDA's dashboards were said to "analyze trends in public opinion and thereby measure the effectiveness" of the campaign, and the affidavit records a quota of 60,000 comments per month for France and Germany combined and a goal to "eventually reach one million" subscribers within a year (paras 61, 66, 73) [129]. China's 912 group rewards members who operate multiple personas "without detection", so the one reward criterion on record measures concealment [135]. Nothing in the texts read shows an operator measuring belief change.

Defenders mostly measure spread or authenticity. The Breakout Scale places an operation in one of six categories by platform span and community span and estimates spread without testing belief change [151]. The EEAS finds that most operations "fail to generate meaningful organic engagement" and that Storm-1516 is "the only IMS capable of generating organic engagement"; most platforms restrict data access, which limits any estimate of magnitude [116]. NATO doctrine says cognitive effect "is not directly observable" [127].

Outcome studies are few and narrow. Bail et al. (2020) surveyed 1,239 US Twitter users in October and November 2017; 76 interacted with an IRA account, and the authors "were not able to find any evidence that interaction with trolls shaped a range of political attitudes and behaviors", from one platform, one country and one month [152]. Eady et al. (2023) found no meaningful relationship between potential exposure to Russian accounts and issue positions, polarization or vote choice in a 1,496-person panel [153]. Of five outcome studies on the commissioning list, three were verified and two were not found; none tests influence operations (Annex C).

Assessment

"Influence alone decides" and "influence has no effect" are both unsupported by the studies reviewed. The record supports a conditional statement: influence can matter in some settings, and the evidence for it measures outputs far more often than outcomes. A claim about effect should say which of the three levels it concerns.

5.7 Reaction to pressure and externalities

The record shows persistence and substitution after takedowns and sanctions. SDA documents call for "multiple 'perishable' accounts", many leased domains were taken down or not renewed, and some moved to other registrars after the designations of 20 March 2024 [129]. That phrase mixes enforcement with ordinary cost decisions, so not every lost domain evidences adaptation. The EEAS reports that Doppelgänger "appears to have discontinued" a practice on Meta platforms, that RRN "experienced a second rebranding" and that "the same FIMI playbook that was deployed in Moldova during 2025" was redeployed in Armenia [116].

One peer-reviewed study measured substitution. Santos Okholm, Ebrahimi Fard and ten Thij (2024) followed 492 public Facebook groups in nine EU countries for three months before and after the 2 March 2022 geo-block of RT and Sputnik. Sharing of their content "fell sharply", non-banned pro-Russian outlets were "largely unaffected", and pro-Russian channels in the Telegram top 50 doubled from 7 to 14 [155]. The study covers public Facebook groups only, counts content and not users, and cannot exclude circumvention.

No source shows an operator stopping because of an indictment or a sanction. One commentator argued that "the idea that these indictments have a deterrent effect strains credulity" (an opinion piece, Tier 3), and Dylan and Maguire (2022) judge that the US and UK disclosure of Russian invasion plans "failed to achieve strategic deterrence", though it may have met narrower aims [154, 156]. The measured cost that appears falls on the defender: on 16 March 2020 the DOJ moved to drop the case against Concord Management, saying a trial would likely expose sensitive law enforcement tools and techniques, "potentially undermining their effectiveness" [157].

Destructive operations spilled past their targets. The October 2020 indictment of six Unit 74455 officers says NotPetya was designed to "masquerade as ransomware" and records nearly USD 1 billion in losses to the three victims identified in the indictment, taken together (paragraph 3(c)) [144]. The Kremlin said it "strongly rejects" the accusations [158]. The Council attributed the Viasat KA-SAT attack to Russia and

placed it "one hour before" the 24 February 2022 invasion; the US presented the spillover as a consequence of an attack meant "to disrupt Ukrainian command and control", and no Russian reply was found [143, 159, 160]. Olympic Destroyer, the documented attempt at a planted artifact, was attributed two years and eight months later, and Peskov called the accusations "blatant Russophobia" [56, 161].

Stat box

Spillover of destructive operations: nearly USD 1 billion in losses to the three victims identified in the October 2020 indictment, taken together (a floor for those victims and no total for NotPetya), and remote monitoring of 5,800 wind turbines in Germany lost for weeks after the Viasat attack. Both are attributions by the cited authorities, and Russia rejects the NotPetya accusation [144, 158, 162].

5.8 A diagnostic matrix and the attribution posture

Table 15 turns the section into questions an analyst can ask of an incident, with the evidence to look for and the tier needed before a claim is stated as a finding. It builds on the three attribution tests of Section 1: technical traceability, multi-source convergence and evidence of intent.

Table 15. Diagnostic matrix for a suspected deniable operation

Question	Evidence to look for	Tier needed
Who acted, and for whom?	Indictment, sanctions notice or converging assessments naming principal and executor	2
Is the claimed identity genuine?	Fit with the named group's history; artifacts conflicting with build or infrastructure evidence	2
Does technical tracing hold up?	Infrastructure and tooling examined by independent teams	2
Do independent sources converge?	Government, private-sector and judicial records agreeing, with no single source repeated	2
Is there evidence of intent to mislead?	Planning documents, pre-incident messaging or a borrowed persona; otherwise label "alleged"	2
Which impact level is claimed?	Output, outcome or strategic effect, each with its own evidence	2, 3
What did the accused say?	Acknowledgment, denial or silence, quoted with date	2

Source: Authors' proposal built on [8, 42, 116, 119, 120]. Tier 3 as a framework; each question asks for Tier 2 evidence before a finding.

No framework reviewed has a slot for deliberate misattribution. The authors therefore propose an untested variable, attribution posture, with four values coded from the public record (Tier 3). Acknowledged: the operator or its sponsoring state has publicly admitted the act. Falsified: an attributing authority finds that the operator planted false indicators to blame a third party. Denied: the accused sponsor rejects a named attribution, and no authority finds planted indicators. Gray: the act runs through a persona, proxy or contractor, an authority links it to a state, and the sponsor makes no specific reply. A case takes the first value that fits, in that order. The record supplies one illustration of each, always as an attribution by the cited authority and with the accused party's reply: Prigozhin's acknowledgment of 7 November 2022 (acknowledged); the Olympic Destroyer artifact and the Proud Boys persona (falsified in the assessments of the UK and the US, which Russia and Iran reject); the Kremlin's rejection of the NotPetya attribution (denied); and the IMCR persona, which the UK Security Minister linked to Quds Force members on 13 July 2026 and on which Iran gave no immediate comment (gray; Section 6). The variable has not been coded for the 38 incidents of Annex B.1 and is illustrative only. It records what an operator says and does, as the cited sources report it, and settles nothing about whether an accusation is correct. Section 6 turns from the model to the documented record.

6. Verified practice: documented cases and actors

Section 5 built a model of the deniable operation. This section tests it against the record. Two limits apply to every row. Only operations that were detected can be observed, and nearly all were also publicly attributed by some authority, so the cases are a sample of what failed to stay hidden. And nearly every attribution is a government's, a prosecutor's or a court's, so it is given as that authority's attribution, with the accused party's reply where the research found one. The identifiers in the first column of each table, such as A1 and I07, are those of the register in Annex B.

6.1 Historical anchors: what a documented false flag looks like

Seven cases from 1931 to 1964 set the benchmark. Archives, courts or memoirs eventually opened each, and Table 16 grades each characterization on what the sources read support.

Table 16. Historical anchors, 1931 to 1964

Case	What the record shows	Principal evidence	Exposure	Grade; tier
A1 Mukden, 1931	Officers staged a rail explosion blamed on Chinese troops	Tokyo Tribunal judgment	Court finding after 1945	Documented; 2
A2 Reichstag fire, 1933	(i) Nazi authorship disputed; (ii) the decree of 28 February exploited the fire	Historians divide; decree text	Unresolved	(i) Alleged; 3. (ii) Documented; 2
A3 Gleiwitz, 1939	SD team staged a Polish attack on a radio station	Abwehr witness at Nuremberg; Hitler's address of 22 August 1939	Six years	Documented; 2
A4 Mainila, 1939	Soviet shelling of its own side, blamed on Finland	Finnish observation; Soviet reports and Khrushchev memoir, via secondary sources	Decades; no formal admission located	Documented; 2
A6 Lavon Affair, 1954	Israeli network bombed Western targets in Egypt, to be blamed on others	Egyptian trial; Israeli inquiries; recognition of members, 2005	Unraveled at once; chain of command disputed	Documented; 2
A8 Northwoods, 1962	Planned only: Joint Chiefs memorandum listing fabricated incidents	Memorandum of 13 March 1962; Kennedy rejected it	Surfaced 1997; released 2001	Documented; 1
A9 Tonkin, 1964	(i) The "second attack" of 4 August did not occur; (ii) fabrication not shown	NSA historian's study; McNamara's 2003 acknowledgment	Study 2001; release 2005	(i) Documented; 2. (ii) Alleged; 3

Source: Authors' compilation from [54, 55, 163, 164, 165, 166, 167, 168, 169, 170, 171, 172, 433]. Mainila and Lavon rest on secondary sources: a broadcaster, a museum board and an EU disinformation report for Mainila, and two press articles and an essay for Lavon. A1, A3, A8 and A9(i) rest on a judgment, transcripts, a memorandum and an official study, and A2(ii) on the decree text (reading routes in Annex C.1.3). Reichstag authorship is a live historiographical dispute, graded on evidence and not on plausibility. A5 (the CIA's Iran operation of 1953) and A7 (the Protocol of Sèvres) are boundary cases that appear in Annex B only (B.1.4), and A5 stays Alleged until the CIA history is read at source.

Four features recur. Evidence of design came from inside the apparatus: a witness relaying what an Abwehr chief told him, a memoir, a tribunal. Confirmation of design took six years at the fastest (Gleiwitz) and decades at the slowest (Mainila); the Lavon network was caught at once, but its chain of command stayed disputed (Table 16). The absence of such evidence about a 2026 incident is therefore the expected state and shows nothing about whether one occurred (*Inference*, Tier 3). Grades attach to characterizations and not to events: Tonkin is Tier 2 as a misreported incident and Tier

3 as a fabrication, because the NSA historian found an effort to make signals intelligence fit the claim and no "smoking gun" of deliberate fabrication [55]. And the documented cases are few. A fact-checking review draws a similar line, counting Gleiwitz, Mukden and Mainila as real, Northwoods as a rejected proposal, and the Maine and Tonkin as exaggerations of real events (corroboration only) [173]. Ruses at sea, such as the Kormoran's Dutch colors in 1941, sit on the lawful side of the line drawn in Section 3 and appear in Annex B as boundary cases.

6.2 Cold War Europe: Italy and Turkey, graded with care

The Italian record merges three layers: what courts found about specific attacks, what commissions concluded or left open, and what authors infer about state or foreign direction. Table 17 keeps them apart. Most court findings are known only through press and legal summaries, docket numbers are unverified, and every grade is provisional until the judgments are read.

Table 17. Italy and Turkey: findings kept apart from inference

Case	Courts and commissions found	Not established	Grade; tier
B1 Piazza Fontana, Milan, 12 December 1969	Anarchists' acquittals final in 1987; far-right Venetian cell implicated; misdirection by officials	A command chain from any state office to the bombers	(i) Attack and misdirection: Contested; 2. (ii) State direction: Alleged; 3
B2 Peteano, 31 May 1972	Venice assizes, 1987: life for two Ordine Nuovo members; three officers convicted of concealment	A Gladio origin of the explosive	Documented; 2
B3 Gladio network, disclosed 24 October 1990	Prime Minister's disclosure; Senate commission reports; European Parliament resolution	Any role in attacks; ten trials ended without convictions	(i) Existence: Documented; 2. (ii) Role in attacks: Alleged; 3
B4 Bologna station, 2 August 1980	Neo-fascists convicted; P2 chief and others convicted of aggravated false accusation	Who ordered it; alleged masterminds died untried	(i) Perpetrators: Documented; 2. (ii) Financing chain: Contested; 2
B5 Istanbul pogrom, 6 and 7 September 1955	Greek court sentence; Turkish trials of 1960 and 1961; a general's later statements	Formal Turkish acknowledgment; any CIA role	Contested; 2

Source: Authors' compilation from [174, 175, 176, 177, 178, 179, 180, 181, 182, 183, 184]. Court findings are known through press and legal summaries; docket numbers unverified.

At Piazza Fontana, Peteano and Bologna courts found far-right perpetrators and misdirection by officials, and at Bologna a P2 chief and others were convicted of false

accusation to shield the perpetrators [181]. That is post-attack misdirection, a neighbor of the false flag, and no scapegoat chosen by design. No court has established a command chain from a state office or NATO to the bombers. The Senate commission's working draft of 9 January 2001, an internal draft and no adopted report, treats institutional responsibility for the misdirection as certain and states that its findings did not confirm the suspicion that stay-behind played a role in the strategy of tension [175]. Two secondary accounts of the 2005 Cassation judgment on Piazza Fontana conflict: one says the Venetian group's part in the massacre was "not proven", the other that the Court affirmed its responsibility [174, 185]. Neither judgment was read, and neither version is treated as settled.

Turkey offers a stronger case of a staged incident exploited to mobilize violence. A bomb at the Turkish consulate in Thessaloniki preceded the riots, and the retired general who organized the counter-guerrilla later called them "an excellent special warfare operation and it reached its goal" [183]. They damaged 5,317 properties, 73 of them churches, and deaths are estimated at 11 to 37 depending on the source [183, 184]. Turkey has not formally acknowledged state responsibility, and the sources treat a CIA link as unconfirmed.

6.3 Attributed operations by actor

The contemporary record differs from the historical one in speed and kind. Attribution arrives within months or years, from governments, prosecutors and courts, and most operations are deniable proxy actions, with staged blame the exception. Table 18 samples eleven of the 38 incidents of the register (2018 to 2026), with the attributing authority, the status on 11 October 2026 and the authors' grades on the three attribution tests of Section 1.

Table 18. Eleven attributed incidents drawn from the register, with the three attribution tests

Incident	Attributed to, and by whom	Tier	T/C/I	Status, 11 October 2026
I07 Leyton arson, UK, 20 March 2024	Wagner-linked network; UK court and police	2	Y/P/Y	Convicted; 7 to 17 years in custody
I15 Attempted arsons in four countries, 2024	GRU-linked group in Russia; Lithuanian prosecutors	2	P/P/Y	Six charged, 16 January 2026
I09 Marywilaska fire, Warsaw, 12 May 2024	Russian services; Polish ministers, 11 May 2025	2	P/P/P	Trial outcome not retrieved

Incident	Attributed to, and by whom	Tier	T/C/I	Status, 11 October 2026
I25 Drones over Poland, 10 September 2025	Russia, "deliberately"; Polish government. A Russian analyst cites a Ukrainian "false flag"	2 / 3	N/P/P	NATO Article 4 invoked; no forensics published
I26 Copenhagen airport closure, 22 September 2025	No actor; Danish police	2	N/N/N	Closed 25 June 2026; drones not confirmed
I06 Volt Typhoon, advisory of 7 February 2024	PRC state-sponsored actors; CISA, NSA, FBI and partners	2	Y/Y/P	A Chinese body calls it a US fabrication
I32 Fake accounts targeting Taiwan, 2025	PRC organs; Taiwan's National Security Bureau	3	R/P/R	Evidence not described in the report relayed
I18 Zeraati stabbing, London, 2024	Iranian state; Old Bailey, 3 July 2026	2	Y/P/Y	Two men convicted; 12 and 8 years
I38 HAYI and IMCR attacks, from March 2026	IRGC Quds Force, "almost certainly"; UK minister, 13 July 2026	2	U/P/P	IMCR banned; Iran gave no immediate comment
I05 North Korean remote IT-worker scheme	DPRK state; DOJ, OFAC and an 11-country alert	2	Y/Y/Y	Ongoing; Spain named once without designee
I10 NoName057(16) arrests, Spain, 20 July 2024	Hacktivists; CCN-CERT calls state backing "probable"	2 / 3	P/P/N	Investigation continues

Source: Authors' compilation from [186, 187, 188, 189, 190, 191, 192, 193, 194, 195, 196, 197, 198]. T/C/I = technical traceability, convergence, intent; Y met, P partly, N not met, U not public, R not retrieved by this review. Grades are the authors' judgment, and no source publishes scores; the full register is in Annex B.1.

Key finding

Across the 38 contemporary incidents of the register (2018 to 2026, Annex B.1), counting a test as met only when it is graded Y, one meets all three tests on the public record, the North Korean IT-worker scheme, and six meet at least two, each of them reaching a court, a seizure or a multi-state technical record. Twenty meet none. If a partly met test also counts, 17 meet all three and four meet none, so the headline moves with the line between Y and P. The four that meet none even partly are of two kinds. No hostile act is established in the Copenhagen closure and the Iberian blackout, and the attribution rests on little in a vendor report of an AI-run espionage campaign and in the Hormuz claims. Every grade is one reader's judgment, six cells await a source that was not retrieved, and a blind second coding is planned (Annex C.1.2).

Russia. Two datasets bracket the sabotage trend under different rules. CSIS counts physical-effect attacks in Europe at 3 in 2022, 12 in 2023 and 34 in 2024, admitting an

incident only with three credible sources pointing to Russian government involvement [199]. ICCT and GLOBSEC count 151 incidents from February 2022 to February 2026 and 172 people, about 95 percent ordinary citizens with no formal intelligence link, and warn that arrests are not findings of guilt [200]. The totals are not comparable. Courts have reached the recruits, and police called Leyton "a clear example of an organisation linked to the Russian state using 'proxies'", a police characterization [186]. The link to a named unit sits mostly in prosecutors' statements, such as Lithuania's claim that orders came from people "connected with Russia's GRU" at EUR 5,000 to 10,000 per action [187]. *Inference*. "Convicted" and "state-directed" are separate findings (Tier 3).

For drones and jamming, the label does a different job. On 10 September 2025 Poland's deputy prime minister said Russian-made drones had breached Polish airspace 19 times and "deliberately targeted" Poland; a Christian Science Monitor report cites at least 14. Russia's defense ministry said it had "no plans" to attack targets in Poland, and an analyst quoted there called the incursion an "operation under a false flag" by Ukrainian special forces [189, 190]. *Inference*. The label appeared on the day, before any forensic publication, and the rebuttals attacked factual premises such as range and footage and left intent aside (Tier 3) [201]. Copenhagen shows the opposite risk: police said on 25 June 2026 that they could not confirm drone activity and closed the case without suspects [191]. Lithuania logged a record 1,022 pilot-reported jamming incidents in June 2025, and no source states intent [202]. Section 4 covers Eagle S. In the sources retrieved, no court or prosecutor alleges that a 2024 to 2026 Russian-attributed incident was staged to frame a third party (Tier 3, an absence in the sources read).

China. Cyber attributions publish technical detail and converge across governments, as with the Volt Typhoon advisory and the 27 August 2025 Salt Typhoon advisory of 13 states, Spain among them [192, 203]. Attributions of influence operations rest on single official claims or vendor reports with little public traceability: Taiwan's National Security Bureau reported more than 45,000 fake accounts and 2.314 million items in 2025, in a relayed account that describes no evidence and comes from a party to the dispute [193]. Counter-attribution has the same structure. China's CVERC called Volt Typhoon "a political farce written, directed, and acted by the U.S. federal government", and the source reporting Harbin police's April 2025 naming of three NSA staff records no US response [204, 205]. *Inference*. Symmetrical standards give such counter-attributions the same Tier 3 label (Tier 3).

Iran. The pattern is persona first. The FBI seized the domains of Handala, which the Department of Justice is reported to have described as a fictitious MOIS identity, on 19 March 2026; on 27 March the persona claimed a breach of the FBI director's email, and the FBI said the data were "historical in nature" [206, 207]. The Irish Times counted 17 low-level attacks in four European countries since early March 2026, claimed by HAYI; an analyst said "We don't have any absolute evidence that Iran is involved", while France's interior minister saw a "direct link" and "subcontractors, often common criminals" [208]. On 13 July the UK Security Minister said Quds Force members had "almost certainly directed" attacks claimed by IMCR, which was banned, and on 24 September Belgium's CUTA head said the Liège perpetrators "certainly" acted on Tehran's orders. The evidence for both is not public, and Iran gave no immediate comment on the UK statement [195, 209, 210]. The lag between persona claim and attribution, four months for IMCR, is the pattern. The Zeraati verdict is the exception: the judge was sure the attack was for the benefit of a foreign power and said the evidence "overwhelmingly points" to the Iranian government [194].

North Korea. The remote IT-worker scheme is the best-documented identity deception in the set. By this manual's taxonomy it neighbors the false flag and is no instance of it, since the deception targets employers' due diligence and makes no third party look responsible for an act. OFAC's designations of 12 March 2026 describe "fraudulent documentation, stolen identities, and fabricated personas"; the Department of Justice announced on 30 June 2025 searches of 29 laptop farms in 16 states; the multilateral monitoring team, in a report of 22 October 2025 that was not read and is relayed by a trade publication, put crypto-asset thefts by North Korean cyber actors at no less than USD 2.84 billion over 21 months [211, 212, 213]. That figure measures theft and not the revenue of the IT-worker scheme, and whether the report links the two was not checked. An 11-country joint alert was issued on 31 July 2026 by authorities of the United States, Japan, South Korea, Australia, Canada, France, Germany, Italy, the Netherlands, New Zealand and the United Kingdom; the page read names the countries and no agency [196]. A companion ISDO report examines the exposure of Spanish organizations [214]. Spain appears once in the research record, as an unassigned facilitator location in OFAC's release, and is not among the alert's signatories.

Non-state actors and AI. The "hacktivist" label works as a deniability layer. Spain's Interior Ministry reported three NoName057(16) arrests on 20 July 2024 and states no link to any state; Europol's Operation Eastwood (16 July 2025) described more than 4,000 supporters rewarded with cryptocurrency and digital badges, and sources differ

on arrests [197, 215, 216]. CCN-CERT writes that the group probably has the backing of more powerful actors; no public evidence for that was found (Tier 3 for state backing) [198]. For AI-enabled deception, the EEAS counted 540 incidents in 2025, 27 percent of them (147 cases, up from 41) involving AI-related techniques, a count that measures detection practice as much as adversary use [116]. No case retrieved shows AI fabrication producing a successful false provocation. The better-supported effects are added doubt, the liar's dividend of Section 7, and added analyst load. A vendor's report of an AI-run Chinese espionage campaign was criticized for publishing no indicators of compromise, and the vendor builds AI models, so it is an interested source (Tier 3) [217, 218].

6.4 Masked identity in cyberspace

Five cyber cases show how identity is borrowed or invented. None has been tried, the accused individuals are at large, and the accused states deny or ignore the findings, so no entry in Table 19 rises above "strongly evidenced but contested" except the Turla advisory.

Table 19. Masked or borrowed identity in five cyber cases

Case	Mask or borrowed identity	Attribution by authority	Grade; tier
D1 TV5Monde, 8 April 2015	"Cyber Caliphate" persona with an Islamic State identity	France, 29 April 2025: APT28, "used by" the GRU; persona not mentioned	Contested; 2
D2 DNC intrusion, 2016	DCLeaks as "American hacktivists"; Guccifer 2.0 as a "lone Romanian hacker"	US grand jury, 13 July 2018: 12 GRU officers; defendants at large	Contested; 2
D3 Olympic Destroyer, February 2018	Planted artifact matching a Lazarus-linked sample	UK, 19 October 2020: Unit 74455, "almost certainly (95%+)"; US indictment	Contested; 2
D4 Turla and OilRig, disclosed October 2019	Use of an Iranian actor's tools and infrastructure	NSA and NCSC, 21 October 2019	Documented; 2
D5 "CyberAv3ngers", from November 2023	Hacktivist persona; some claims false	CISA, 1 December 2023: IRGC-affiliated; Treasury sanctions, 2 February 2024	Contested; 2

Source: Authors' compilation from [56, 60, 61, 62, 219, 220, 221, 222, 223, 224].

Of the five, only Olympic Destroyer is admitted by Proposal A of Section 3; TV5Monde is undecided between a false flag and an opportunistic cover. A vendor found a copied metadata artifact that conflicted with the binary's own build characteristics and called

the pattern an "intricate false flag operation" without naming a group; the UK said the unit "attempted to disguise itself as North Korean and Chinese hackers", and the US release as read has no false-flag language [56, 60, 61]. The DNC personas masked the actor without blaming a third party. The indictment alleges that the defendants searched their Moscow server for English words that later appeared in the first post of a "lone Romanian hacker", and the Kremlin spokesman denied involvement on 18 July 2016 [221, 222]. Turla is the cleanest borrowing, and the advisory judges that the Iranian tools' operators were "almost certainly not aware of, or complicit with" the use [62]. Vendors read the pattern within weeks, FireEye on TV5Monde in April 2015 and Kaspersky on Olympic Destroyer in March 2018; governments attributed it ten years and 2 years 8 months later [60, 220]. *Inference*. In cyberspace a persona or planted artifact is a design choice, and the attribution-posture variable of Section 5 would record it as "falsified" only after an authority has said so (Tier 3).

6.5 State operations under denial

Four cases show what follows a state's denial of a deniable operation. In Crimea the public denial lasted six weeks: Putin said on 4 March 2014 that the armed men were "local self-defense forces" and on 17 April that "the Russian servicemen did back the Crimean self-defence forces" [225, 226]. Counted from 27 February 2014, the date from which the Strasbourg Grand Chamber found Russian effective control, the interval from event to admission is seven weeks [227]. On MH17 the Hague District Court convicted three men in absentia on 17 November 2022, and the Kremlin spokesman called the Strasbourg and Hague rulings "null and void" (Section 4) [228]. In Salisbury, charges came six months after the attack and EU listings of GRU officers about ten and a half months after it. Summaries of the public inquiry's report of 4 December 2025 say it found that GRU members carried out the attack and that the President authorized it, a finding read only through those summaries. Russian officials reject the attribution, and the research found no report of a trial [229, 230, 231]. Nord Stream runs the other way: a German indictment announced on 2 July 2026 alleges that Ukrainian state authorities ordered the sabotage, which Ukraine denies, and the trial was due to open on 14 October 2026, three days after the evidence cutoff. The allegation stays at Tier 3 until a verdict [232, 233]. Update of October 22, 2026, after the evidence cutoff: the trial opened on 14 October 2026 in the Hamburg Higher Regional Court, the defendant denies any involvement, and the trial is expected to run into February 2027 [428]. Legal-press reports count 16 hearing dates to 17 December 2026 and more planned for January and February 2027 [233, 427]. The court's own schedule was not

read, and nothing in these reports changes the Tier 3 status of the allegation of state direction.

Inference. The cases are selected by the existence of a record, so they show how long denial lasted where it ended and nothing about operations where it did not. Attribution arrives in layers, and two intervals, both counted from the event, separate them. The first ends at the earliest on-the-record statement by a government that names a state, and anonymous briefings and company reports do not count. The second ends at a court finding, an official inquiry report or an admission by the accused state. The first public attribution by a government that the research located came at most 2 days after the event in Crimea (Ukraine's statement to the Security Council on 1 March 2014) [430], 8 days after it in Salisbury (the Prime Minister on 12 March 2018) [431], 2 years 8 months after it for Olympic Destroyer (19 October 2020) [56] and about 10 years after it for TV5Monde (29 April 2025) [219]. A final finding or admission came about seven weeks after the event in Crimea (Putin, 17 April 2014) [226] and about 7 years 9 months after it in Salisbury (the inquiry report of 4 December 2025, read through summaries) [231]; no court finding, inquiry report or admission was located for Olympic Destroyer or TV5Monde. In Salisbury, charges at six months and EU listings at about 10.5 months fell between the two (Tier 3). Figure 4 plots the intervals. Section 7 turns from what the record supports to what it does not.




Figure 4. Time from the event to the first public attribution by a government and to a final finding or admission, four cases (logarithmic scale).

Source: [56, 219, 226, 227, 229, 230, 231, 430, 431]. Days are counted from the event: Crimea 27 February 2014, Salisbury 4 March 2018 [431], Olympic Destroyer 9 February 2018 and TV5Monde 8 April 2015. The dotted line runs to the cutoff of 11 October 2026. The exact dates are in the open data (attribution_intervals.csv).

7. Unconfirmed narratives and the misuse of the label

Section 6 listed what the record supports. The label also travels well beyond that record. It is attached to mass shootings, bridge collapses, wartime strikes and election results, usually by someone who needs an event to belong to somebody else. This section grades claims and leaves the events alone. One rule from Section 1 governs it: "Refuted" is used only where an institution or court examined the specific claim, and every other claim stays "Alleged". In practice a claim is Refuted when the examination reached the facts the claim denies at first hand and found a different cause, or no basis for the claim at the sites or events examined (Bucha, Sandy Hook, the Key Bridge). It stays Alleged when no body examined it, when the body declined to attribute, when only an interested party examined it, or when the examination supported part of the claim, as the Italian courts did in finding official misdirection (Gladio). All rows are Tier 3.

7.1 Four families, graded claim by claim

The families are reciprocal accusations in the Ukraine war (a), claims that the attacked state organized the incident itself (b), attacks, disasters and health crises presented as staged (c), and the claim that psychological warfare alone explains shifts in opinion (d). Table 20 sets fifteen claims against the institutional record. The authors chose the claims within the four families and applied no formal selection rule, so the set is a selection and no sample, and eight of fifteen is no rate of false accusations.

Table 20. Fifteen claims graded against the institutional record

ID, family	Claim examined	What institutions found	Grade; tier
N1, a	Bucha killings were staged	OHCHR documented 73 civilian killings; Maxar imagery appears to show bodies under Russian control	Refuted; 3
N2, a	Ukraine prepared a "dirty bomb"	IAEA, at Ukraine's request, found "no indications of undeclared nuclear activities" at three sites	Refuted at those sites; 3
N3, a	A HIMARS strike caused the Olenivka explosions	OHCHR: not caused by HIMARS rockets; weapon and origin undetermined, no access	Refuted as to weapon; 3
N4, a	Kakhovka dam and Zaporizhzhia plant attacks blamed on the other side	UN commission made no attribution; IAEA declines attribution by design	Alleged; 3

ID, family	Claim examined	What institutions found	Grade; tier
N5, a and b	Each side's account of the 3 May 2023 Kremlin drone attack	No institutional finding; analysts and anonymous US officials, as relayed, disagree	Alleged; 3
N6, b	Britain staged the Salisbury poisoning	OPCW answered the "BZ" point as a control-sample issue; EU listing names GRU officers	Refuted; 3
N7, b	The Pahalgam attack of 22 April 2025 was a false flag	India's chargesheet names Lashkar-e-Taiba and an affiliate; no independent investigation located	Alleged; 3
N8, b	Gladio, the CIA or NATO directed the Italian bombings	Courts found far-right perpetrators and official misdirection, but no state or NATO command chain	Alleged; 3
N9, c	The 9/11 attacks and the WTC 7 collapse were staged	Commission: 19 attackers for Islamist extremists; NIST: no evidence of a blast at WTC 7	Refuted; 3
N10, c	Sandy Hook was a staged hoax	State's Attorney: shooter "acted alone"; no co-conspirators found	Refuted; 3
N11, c	The 2024 Key Bridge collapse was an attack	FBI, 27 March 2024: no specific, credible information of terrorism ties; NTSB: a loose wire started two blackouts	Refuted; 3
N12, c	Directed-energy weapons started the 2023 Lahaina fire	Maui Fire Department and ATF: accidental, from broken power lines; afternoon ignition not definitive	Refuted; 3
N13a, c	COVID-19 was a staged hoax	No institutional finding on a staged hoax in the sources read; ODNi and WHO address origin	Alleged; 3
N13b, c	COVID-19 was a deliberate release	ODNi: not developed as a bioweapon, a narrower claim; WHO 2025: evidence favors spillover, laboratory leak not excluded	Alleged; 3
N14, d	Russian influence alone decided the 2016 US vote	ODNi made no impact assessment; a 1,496-person panel found no meaningful link to vote choice	Alleged; 3

Source: Authors' compilation from [153, 174, 175, 181, 230, 234, 235, 236, 237, 238, 239, 240, 241, 242, 243, 244, 245, 246, 247, 248, 249, 250, 251, 252, 253, 432]. "Refuted" means an institution or court examined the specific claim. N2 holds only for the three sites inspected, with sample results pending when the IAEA spoke. N3 holds as to the weapon; responsibility for Olenivka is open. N6 rests on an AP report of the OPCW statement, N9 on the Commission's summary and NIST's FAQ, and N12 on a news report of the official finding. N1 and N10 rest on examinations of the killings that the claims deny. N11 rests on the NTSB's cause and on the FBI statement of 27 March 2024, which reported no specific and credible information to suggest ties to terrorism "at this time" [432]. India's attribution in N7 is also untested in court.

Key finding

Eight of the fifteen claims are graded Refuted, and in each an inquiry, court or inspection body had examined the specific claim. The other seven stay Alleged. That grade says a claim is unproven and does not say it is false. The fifteen are a selected set within four families, so the split is no rate of false accusations.

The split between the grades follows access. At Bucha, OHCHR made site visits and interviews from 9 April 2022, after Russian forces had left, and the question was settled [234]. Where access was denied or the mandate excluded attribution, the accusations stayed open. At Olenivka, OHCHR excluded the HIMARS explanation from witness interviews and photo and video analysis but could not name the weapon: Russia gave no UN access, and the scene "was not preserved but was instead contaminated" [238]. At Zaporizhzhia, the director general of the IAEA said on 19 April 2024 that the proof of an attack's origin his agency would need was "simply impossible" to obtain, since a drone "could be manned from one side and belong to another side" [239]. *Inference.* A detection framework should log the access conditions of an incident before it weighs either side's statements (Tier 3).

Items outside the four families show the same structure. Iranian officials suggested a US staging of the 2019 Gulf of Oman tanker attacks, while the United Kingdom called Iranian responsibility "almost certain"; no independent forensic finding appears in the sources read [254, 255]. In 2026 Iran's foreign ministry spokesman spoke of possible "false flag" operations after Turkey reported an intercepted missile, and called the Diego Garcia launch an "Israeli false flag"; on the day of Flydubai flight FZ1073, 30 September, security officials said Israel meant to stage an attack and blame Iran. No evidence accompanied any of the three statements, and by 3 October the UAE attorney general had called the co-pilot's act a "terrorist act" without naming any Iranian role [256, 257, 258, 259, 260].

On Douma, Russia's chief of the general staff said on 13 March 2018, 25 days before the attack, that rebels and the White Helmets planned a staged chemical attack. The OPCW Fact-Finding Mission found reasonable grounds that a toxic chemical, likely chlorine, was used, and its Investigation and Identification Team concluded in January 2023 that there were reasonable grounds to believe Syrian Arab Air Forces were the perpetrators [261, 262, 263]. The strongest technical dissent, a leaked staff memo of 27 February 2019, argued that both cylinders were more probably placed by hand. That concerns placement, and no inquiry has established staging. The Nord Stream account naming the United States and Norway rests on one unnamed source in a

newsletter of 8 February 2023, which the White House called "false, a complete fiction"; Section 6 covers the German case [264].

7.2 What the label does in argument

The label does three jobs in these rows, and each can be seen without judging anyone's beliefs. It answers an accusation by naming the accuser as the author: Russian officials called the Bucha footage a "hoax" and a "staged production" [236]. It pre-empts one, as in the Flydubai statements and the Douma announcement above. And it explains an opponent's opinion as enemy influence, which removes any need to engage with the opinion. *Inference*. The third use is the narrative weapon of the fourth family (Tier 3).

The evidence on that family points away from "alone". The ODNI statement of 6 January 2017 says the intelligence community "did not make an assessment of the impact" of Russian activities on the 2016 outcome [253]. A three-wave panel of 1,496 US respondents, matched to the Twitter feeds they followed, found that 70 percent had at least one potential exposure to a foreign influence campaign and that 1 percent of users accounted for 70 percent of exposures. It found no meaningful relationship with vote choice, an estimated change in the Trump vote of about -0.18 points [153]. The design is observational and covers one platform. Elsewhere exposure did move votes: Russian television raised support for pro-Russian parties in Ukraine in 2014, and independent television cut the government party's vote share in Russia in 1999 by about 8.9 points, both read at abstract level only [265, 266]. No study read supports "alone", and none supports "no effect".

A mirror image has experimental support. Schiff, Schiff and Bueno (2025) ran five pre-registered survey experiments with 15,287 US adults. Politicians who falsely called a real text-based scandal "fake news" or a "deepfake" gained support, by about 0.17 standard deviations for creating uncertainty and 0.21 for rallying supporters, and the tactic was largely ineffective against video [267]. The term "liar's dividend" comes from Chesney and Citron [268]. *Inference*. Saying "it was staged" about a real incident is the same move as saying "it is fake" about a real scandal. No study retrieved tests that analogy for false-flag claims (Tier 3).

Belief research limits what an analyst may infer from who holds a claim. A meta-analysis of 74 effect sizes from 13 studies found no pooled association between any Big Five trait and conspiracy belief [269]. Paranoia and conspiracy belief correlate but are distinct constructs in a meta-analysis of 11 studies plus two further studies (N = 578),

and 26.7 percent of 5,645 respondents in a US national survey endorsed a general conspiracy item [270, 271]. No study read shows that belief in a false-flag narrative is a clinical condition, so this manual uses no clinical language. And because real false flags exist (Section 6), a measure that counts every such belief as irrational would misclassify the ones that are correct for some events.

7.3 A disputed article and the evidence ladder

One disputed secondary source needs separate handling. "NATO/US false-flag attacks in Europe", by David Ray Griffin and Clifford W. Cobb, appeared in the *American Journal of Economics and Sociology* 82(5), November 2023, page 455 [272]. A Wikipedia page describes the journal as peer reviewed, published by Wiley-Blackwell and founded in 1941; the same page lists Cobb as editor-in-chief, which was not checked for 2023 [273]. Griffin (1939 to 2022), a founding member of Scholars for 9/11 Truth, died in November 2022 [274]. The abstract lists no cases or sources, and the full text and references were not read. This manual cites the article for no finding and treats it as the statement of a claim.

The abstract bundles three propositions of different standing, and Section 6 separates them. Established (Tier 2): stay-behind networks existed and were concealed from parliaments, and courts found far-right perpetrators and official misdirection at Piazza Fontana, Peteano and Bologna. Not established (Tier 3): that Gladio, the CIA or NATO directed the attacks. Il Post counts ten Gladio trials with no convictions, and the Senate commission's working draft says its findings did not confirm the suspicion [175, 180]. The literature behind the stronger reading rests on a narrower base: works that reviewers received in opposite ways, and US Army Field Manual 30-31B, which a secondary summary says the CIA called a KGB forgery in 1980 [275, 276]. A structural problem remains. The thesis that false flags are unknowable because secrecy hides them cannot be tested, and an untestable claim has no grade. This manual grades each case on what courts, commissions and archives show.

7.4 The cost of error in both directions

Speculation can run ahead of evidence, and Spain supplies a case. On 29 April 2025, the day after the Iberian blackout, an Audiencia Nacional judge opened proceedings on possible cyberterrorism, the prime minister said no hypothesis was ruled out, and Red Eléctrica's Eduardo Prieto said there was no intrusion of any kind in the control system [277]. The government committee reported on 17 June 2025 that no evidence existed of a cyber incident as the cause [278]. The court published a provisional

dismissal on 12 January 2026, finding no "mínimo indicio" of terrorist sabotage, and ENTSO-E's final report of 20 March 2026 described a combination of many interacting factors [279, 280]. Nobody attributed the blackout to an actor, so it is no case of over-attribution. A hypothesis of hostile action stayed open until the court, the government committee and ENTSO-E each reported no cyber cause (all three attribution tests graded N). Over-attribution proper is better shown by I27, the Brussels airport drones of 4 November 2025, where "most likely Russia" rested on anonymous sources and a minister (graded N, N and P; Moscow denies). The claim outran its evidence without being shown false (Table 34).

Under-attribution costs time. The first public attribution by a government came 2 years 8 months after Olympic Destroyer and about 10 years after TV5Monde, and the inquiry report on Salisbury came about 7 years 9 months after the attack (Section 6.5). An analyst who answers one error with the other loses either way. *Inference*. A "false flag" assertion should be handled as a hypothesis with a stated evidential threshold in both directions, a rule that Section 8 turns into method (Tier 3).

8. Detection, assessment and countermeasures

Section 7 closed on a rule: a claim that an incident was staged and a claim that it was not each need a stated threshold of evidence. This section turns the rule into method. It reviews the standards analysts already use, the indicators open sources support, a procedure with thresholds in both directions, the record of countermeasures, and the exposure of companies and of Spain. Standards and institutional records are Tier 1 or Tier 2 material; the procedure and its thresholds are the authors' proposal (Tier 3). Several primary documents were read through extraction tools, so quotations need rechecking against the originals (Annex C).

8.1 Standards for judgment and their evidence base

Official standards agree on three requirements: alternative hypotheses, calibrated uncertainty language and an explicit check for denial and deception. ICD 203, signed on 2 January 2015, lists nine analytic tradecraft standards. One says that "Analytic products should identify and assess plausible alternative hypotheses", and the sourcing standard names "possible denial and deception" as a factor to weigh [281]. The CIA's 2009 Tradecraft Primer builds the same check into analysis of competing hypotheses (ACH), which asks "Is denial and deception a possibility?" [282]. As extracted, the Primer does not discuss planted evidence explicitly.

The two best-known scales differ (Table 21). ICD 203 strongly encourages analysts not to mix terms from different rows, and it forbids combining a confidence level and a likelihood term in one sentence. The UK keeps confidence in a separate High, Moderate or Low rating built from the information base, analytical rigor, and complexity and volatility [101, 281]. *Inference*. A statement that mixed the two vocabularies would invite misreading, so any attribution statement should name its scale (Tier 3).

Table 21. Two official probability vocabularies, side by side

Row	ICD 203 (United States)	PHIA yardstick (United Kingdom)
1	Almost no chance, remote: 1 to 5 percent	Remote chance: above 0 to about 5 percent
2	Very unlikely, highly improbable: 5 to 20 percent	Highly unlikely: about 10 to 20 percent
3	Unlikely, improbable: 20 to 45 percent	Unlikely: about 25 to 35 percent
4	Roughly even chance, roughly even odds: 45 to 55 percent	Realistic possibility: about 40 to under 50 percent
5	Likely, probable: 55 to 80 percent	Likely or probable: about 55 to 75 percent
6	Very likely, highly probable: 80 to 95 percent	Highly likely: about 80 to 90 percent
7	Almost certain, nearly certain: 95 to 99 percent	Almost certain: about 95 to under 100 percent

Source: Authors' compilation from [101, 281]. The ICD 203 pairing of terms with ranges was checked against the PDF of the directive for the second edition. The PHIA ranges were checked against the GOV.UK publication on October 22, 2026; PHIA stands for Professional Head of Intelligence Assessment.

The evidence that structured techniques raise accuracy is thin. A RAND report released on 9 January 2017 found that the intelligence community had "made little effort to assess whether they are measurably improving the quality of analysis"; its own pilot was qualitative, and the page gives no sample [283]. The one randomized test on practicing analysts that was located, by Dhimi, Belton and Mandel (2019), assigned 50 analysts at a UK intelligence organization, 25 trained in ACH and 25 controls, to judge which of four fictitious tribes a target belonged to. Trained analysts more often ran a sensitivity analysis (60 percent against 4) and considered diagnosticity (80 against 32), but they did not pick the correct hypothesis more often (36 against 33) and used base rates less often (12 against 52) [284]. Checked against the full text on October 22, 2026, the six figures reconcile: the control group's 33 percent is 8 of 24 analysts (N = 49 for that test), and the other percentages are consistent with whole numbers of the 25 analysts in each group. The authors also report lower consistency of evidence assessment under ACH and caution that the technique may increase inconsistency

and error in judgement [284]. That caution bears on Annex D, which asks a second analyst to repeat the key tests blind (D.1.6), so that inconsistency shows as a difference between replications (Tier 3). *Inference.* The techniques are best treated as audit trails and challenge functions, and the weak use of base rates matters for attribution, where the question is how often an actor family stages or imitates another. No study located tests the techniques on attribution tasks (Tier 3).

8.2 Indicators and open-source validation

The Berkeley Protocol on Digital Open Source Investigations (2 December 2020, UC Berkeley Human Rights Center with OHCHR) is the closest standard for verifying open-source material. It separates credibility, the believability of information, from reliability, the dependability of a source (para. 23), warns that creators "can now spread disinformation and manipulate digital data with relative ease" (para. 7), and weighs each item holistically, including whether other information corroborates or contradicts it (para. 57) [285]. It was written for international crimes, so using it on hybrid incidents extends its method. For the second edition the full text was read, including the chapter on geolocation, chronolocation and hashing, and Annex D applies it.

Documented detection successes are few, and each rests on one evidence class. Meduza reported that Bellingcat's Aric Toler verified video metadata showing that the evacuation announcements of the Donetsk and Luhansk leaders, shared on 18 February 2022, had been recorded on 16 February [286]. The article names no metadata fields, and Bellingcat's own write-up was not retrieved, so the check stands as one evidence class awaiting corroboration. The UK statement of 19 October 2020 on a GRU unit posing as North Korean and Chinese hackers describes no technical evidence on the page read, and the NSA and NCSC advisory documents Turla's use of an Iranian actor's tools [56, 62]. None of these comes with a false-positive rate.

Institutions publish criteria of their own. VIGINUM defines foreign digital interference by four criteria: manifestly inaccurate or misleading content; artificial or automated, massive and deliberate dissemination; intent to harm the fundamental interests of the Nation; and direct or indirect involvement of a foreign actor. It works from open sources only and does no fact-checking [287]. CSIS admits an incident to its dataset only with at least three credible sources, assigns each a confidence level, and records that a BAE Systems explosion in South Wales was not confirmed as sabotage [199]. Table 22 collects the classes these sources support.

Table 22. Indicator classes assembled from the sources (authors' synthesis)

Indicator class	Support in sources	Verification route	Main false-positive risk
Timing or provenance mismatch	Meduza account of a Bellingcat check; Berkeley paras 55 to 57	File-property analysis plus independent corroboration of time and place	Metadata can be wrong, edited, reset or re-uploaded
Narrative pre-positioning	EEAS 4th report; VIGINUM short-term mode; US pretext warnings	Cross-platform timeline; coded incident records	Fast but legitimate commentary; self-negating warnings
Signature mismatch	UK 2020 Olympics statement; NCSC and NSA Turla advisory	Independent evidence classes at several layers; cross-agency review	Genuine tool sharing, reuse or leaks
Single-channel evidence	CIA 1972 text; ICD 203 sourcing; Berkeley reliability and credibility	Corroboration through independent channels; source and information rated separately	Real events with restricted access
Impersonation and covert amplification	EEAS 4th report; VIGINUM four criteria	Open-source mapping of channels and infrastructure	Organic coordinated activism
Anomalies in a claim of responsibility	No institutional source retrieved	Not established	Not established

Source: Authors' synthesis, Tier 3, from [56, 62, 116, 281, 285, 286, 287, 288]. No class has a validated error rate.

A timing mismatch discriminates only after the metadata itself is authenticated. Physical evidence, likewise, does not establish intent or an actor by itself: the BAE explosion was not confirmed as sabotage, and Eagle S shows how far prosecutable proof can lag political attribution (Section 4).

8.3 A two-sided procedure for analysts

The procedure below combines the standards above with the tests of Section 1 and the evidence rule of Section 3. It applies to a claim of staging, to the attribution that claim contests, and to the possibility that no hostile act occurred. It is the authors' proposal (Tier 3), untested against cases. Annex D (D.1 to D.3) supplies the routines for evidence, timing and wording that the steps assume. Figure 5 shows the seven steps in sequence.

1. Record the claim and its channel: who said it, when and through which source. Rate the source and the information separately.
2. Log the access conditions: who could inspect the scene, logs or wreckage, and who could not (Section 7).
3. State three hypotheses: H1, the named party acted; H2, the act was staged or disguised to blame another; H3, no hostile act occurred.
4. Write down, for each, the evidence that would count against it. Without this step, deception can explain any gap.
5. Grade T, C and I (Y, P, N, U or R, as in Table 2) for H1 and H2 and give each item its tier; for H3, use the criterion of Table 23 (evidence classes examined and a positive cause identified). Discount evidence that fits all three hypotheses, since motive generates hypotheses and rarely separates them.
6. Apply Table 23, then state the judgment on a named probability scale with a separate confidence level and what would raise it [101, 281].
7. Set a review trigger, such as a court filing or an official report, and record every change of judgment with its reason.




Figure 5. The two-sided procedure of Section 8.3 as a sequence of seven steps (authors' proposal, Tier 3). Source: Authors' proposal, Tier 3, as listed in Section 8.3 and Table 23; it has not been tested against cases. T, C and I are the tests of Section 1. The return arrow shows the review trigger of step 7.

Table 23. Evidential thresholds for three hypotheses (authors' proposal)

Hypothesis	Keep open	Assess as likely or higher	Reject, or grade Refuted
H1: the named party acted	Any attribution on record; source and access conditions logged	At least two of T, C and I graded Y, from independent sources	A tested alternative explains the evidence better, or a core artifact is shown planted
H2: staged or disguised to blame another	One indicator, a third party with motive, and disconfirming evidence named in advance	Tier 2 evidence of design: planning papers, a traced forgery or an official finding	An institution or court examined the specific claim and found no staging
H3: no hostile act	Physical facts fit an ordinary cause; hostile causes not yet tested	Named evidence classes examined, none hostile, and a positive cause identified	Evidence of deliberate interference in any examined class, or an official finding of intent

Source: Authors' proposal, Tier 3, built on [199, 281, 282, 285]. T, C and I are the tests of Section 1. The thresholds map evidence classes to wording and carry no measured error rates.

The thresholds are asymmetric in one place and symmetric in another. Keeping H2 open costs little, and ignoring deception has a documented history: a 1972 CIA text records that analysts accepted the 1968 Czechoslovakia "exercises" at face value because earlier announced exercises had been real [288]. But the bar for asserting H2 equals the bar for asserting H1, Tier 2 evidence, because declaring staging without discriminating evidence recreates the liar's dividend for the declarer (Section 7). The word Refuted stays reserved, under the rule of Section 1, for claims an institution or court has examined.

The manual's own register is the only empirical material to hand, and it cannot validate these thresholds. In the contemporary record staged blame is the exception, but a Y grade on T or I is mostly awarded because an indictment, a seizure or a multi-state advisory exists, so the six incidents that meet at least two tests were bound to have reached one of them (Section 6). The register illustrates the thresholds and does not test them. A test is still missing: grade cases whose outcome is now known (Salisbury, Olympic Destroyer, NotPetya, the DNC intrusion) from the evidence public at an early date, record what Table 23 would have returned and compare it with the later finding. It has not been run. Applied to the register as graded today, Table 23 returns "keep open" for H1 on Salisbury (I01, graded R/P/P) although a statutory public inquiry has since reported, on the summaries read (Section 6.5). That shows how the thresholds and the grades interact, and why the blind recoding of Annex C.1.2 matters. Under-calling has a price as well, since the interval from event to first public attribution ran from at most 2 days to about 10 years, and to a final finding or admission from

about seven weeks to about 7 years 9 months (Section 6.5). The Iberian blackout shows H3 done properly. The government committee reported on 17 June 2025 that it found no indications or evidence of a cyber incident in the entities analyzed, after reviewing IT systems, control-center protection systems, traces of malicious code and network traces, and it noted vulnerabilities with no evidence that they were exploited [289]. These evidence classes come from a press account, because the report itself was not read.

8.4 What the countermeasure record shows

The institutional stack is dense and still growing. The EU Hybrid Toolbox followed the Strategic Compass of March 2022, a framework for Hybrid Rapid Response Teams was approved in May 2024, and the Council conclusions of 16 March 2026 call for "increasing the cost of hybrid activity against the EU for those responsible" [34, 290]. NATO created a Special Coordinator for Hybrid Threats in 2025 [291]. Table 24 records what the sources show about effect.

Table 24. Hybrid countermeasures and the evidence of their effect

Measure	Body and date	Evidence of effect	Limit
Hybrid sanctions regime	Council; first listings 16 December 2024	Output only: 59 individuals, 17 entities by December 2025; 69 and 19 by April 2026	No outcome evaluation found; legal challenges not checked
FIMI Toolbox, Rapid Alert System, threat reports	EEAS; 4th report 12 March 2026	540 incidents in 2025, 65 percent unattributed; ECA 2021: no overall KPIs	Sample is "selective, time-bound"; ECA finding dates from 2021
Counter-hybrid strategy, support teams, Special Coordinator	NATO; 2015, 2018, 2025	No published evaluation found	Summit texts of 2025 and 2026 are short (Section 2)
Public attribution with sanctions	UK Foreign Office and NCSC	Government says the Sturgess Inquiry report of 4 December 2025 confirms its long-held assessment	No deterrence evaluation; Russia rejects the attribution (Section 6)
Foreign digital interference unit	VIGINUM, France; decree of 13 July 2021	77 Storm-1516 operations documented to 5 March 2025	Impact "difficult to estimate"; open sources only

Source: Authors' compilation from [90, 116, 287, 291, 292, 293, 294, 295, 296]. Counts of listings are the Council's cumulative totals for the dates shown; Section 9 returns to later rounds.

Key finding

Two measures outside Table 24 have quantified effect studies, the RT and Sputnik ban (Section 5) and pre-bunking videos (below), and each is narrow. No institutional instrument has an outcome evaluation in the sources read, and the incident counts of the EEAS give no trend.

Inference. Reports count outputs. The EEAS counted 750 incidents between 1 December 2022 and 30 November 2023 and 540 in calendar 2025 from what it calls a "selective, time-bound sample", so the two figures cannot be compared [116, 297]. The European Court of Auditors found in 2021 "No overall KPIs" and metrics that "satisfy only output indicators", a description that still fits the 2026 report [294]. The EU appears to have no published hybrid-attribution procedure, the UK and France publish unit-level or campaign-level attributions, and Spain reserves public attribution to a third State for its National Security Council. All treat attribution as a political decision resting on technical assessments (Tier 3).

8.5 Pre-bunking, disclosure and overreach

Pre-bunking has the strongest experimental base of the countermeasures reviewed, and two later studies qualify it. Roozenbeek and colleagues (2022) tested five short videos on manipulation techniques in six randomized studies ($n = 6,464$) and a YouTube field study ($n = 22,632$); the videos improved recognition of the techniques, discrimination between trustworthy and untrustworthy content and sharing decisions [298]. A reanalysis of five studies of two gamified interventions found that the games did not improve discrimination and made participants answer "false" more often to every item [299]. Three online experiments in the United States, Poland and Hong Kong ($n = 6,127$) found that each intervention reduced belief in false information and also lowered the perceived credibility of factual information [300]. All three were read at abstract or record level, and the first study's field outcome is a short post-video test. *Inference.* Evaluations for hybrid incidents should report discernment and not belief in falsehoods alone, and no evidence on durability over months was retrieved (Tier 3).

Pre-emptive disclosure has costs that its defenders name too. Dylan and Maguire (2022) list sources lost once exposed, narrowed escalation options, a self-negating prophecy in which a warning that prevents an attack looks wrong, and credibility damage when claims lack support [154]. The version read is an accepted manuscript with a different working title, so quotations need checking against the published text.

Overreach is the other risk. In Romania the Constitutional Court annulled the 2024 presidential election process on 6 December 2024 (Decision no. 32, file 3771F/2024) [420], two days after declassified intelligence documents were disclosed. A legal scholar argues that the Court acted on its own motion although the statute allows annulment only on candidates' referrals within three days of voting, and that it left the 1 December parliamentary elections standing despite similar TikTok findings; she also notes that European standards permit annulment when proportionate [301]. This is expert commentary on process, and no independent evaluation of the disclosure was retrieved. The one international judicial step is narrow: on 6 March 2025 a committee of the European Court of Human Rights declared Călin Georgescu's application (no. 37327/24) inadmissible, holding that the presidency is outside Article 3 of Protocol No. 1, and it did not reach the merits of the annulment (Annex B.2) [418, 419]. Csernaton and Wanless (2026) warn of "threat inflation", the blaming of foreign actors for largely domestic problems such as engagement-driven platform design [302]. Some listing grounds in the hybrid sanctions regime are speech-adjacent, such as analysts "closely linked to or affiliated with the Kremlin's policy and messaging apparatus", and the sources report no legal challenge or civil-liberties assessment [292].

No dataset separates confirmed false flags from mislabeled incidents, so no false-positive rate can be given. Safeguards that appear in the material include published criteria, open-source-only data, an independent ethics committee reporting to the Prime Minister (VIGINUM), annual renewal of sanctions, and parliamentary or judicial oversight. Their effect on false-positive rates is untested (Tier 3).

8.6 Companies and critical-infrastructure operators

Section 4 noted that NIS2 reporting attaches to the incident and never to the attacker's identity. Table 25 shows the clocks of Article 23, and Annex D sets them beside the other legal clocks that shape a response (Table 66).

Table 25. NIS2 incident-reporting clocks (Article 23)

Stage	Deadline	What the text requires
Early warning	Within 24 hours of becoming aware	Whether unlawful or malicious acts are suspected; possible cross-border impact
Incident notification	Within 72 hours	Initial severity and impact assessment; indicators of compromise where available
Final report	Not later than one month after the notification	"The type of threat or root cause that is likely to have triggered the incident"

Source: Authors' compilation from [106, 303]. The deadlines were checked against the official EUR-Lex text of Article 23 for the second edition and agree with recital 102.

On the text read, no stage requires naming the attacker. *Inference.* A company statement can therefore be complete at the level NIS2 demands while saying the actor is unknown; this reads Article 23 and is no regulator's interpretation (Tier 3). The stakes are real: maximum fines for essential entities are at least EUR 10 million or 2 percent of worldwide annual turnover, whichever is higher, and for important entities at least EUR 7 million or 1.4 percent (Article 34) [304]. NIST's incident-handling guidance, SP 800-61 Rev. 3 of April 2025, asks response teams only to attempt to determine which threat actors were involved (RS.AN-03) and sets out no attribution method [403]. Rev. 2, withdrawn on 3 April 2025, had called identifying an attacking host "a time-consuming and futile process that can prevent a team from achieving its primary goal" [305]; that sentence does not appear in Rev. 3. Rev. 3 adds that formal chain-of-custody handling "might not be performed for every incident" (RS.AN-07) [403]. Where an organization does keep a chain of custody, its rules are the corporate counterpart of the Berkeley Protocol's preservation rules (Annex D).

Government practice supplies models of calibrated language. The NSA and NCSC advisory on Turla concedes that "attribution of attacks and proving authorship of tools can be very difficult", and the UK's Olympics statement pairs "high confidence" with "almost certainly (95%+)" [56, 62]. *Inference.* A crisis statement can follow five parts: what is known; what is assessed, with a named probability scale and a separate confidence level; what is not known; what would change the assessment; and when the next update will come. It names no actor before evidence supports one and suggests no cyber cause where none has been shown, as in the Iberian blackout (Tier 3).

Identity risk shows why telemetry alone cannot close a case. The US guidance of 16 May 2022 on North Korean IT workers lists red flags that fall to different functions: inconsistent names or work histories (human resources), equipment shipped to an unlisted address (security), payment requests in virtual currency (finance), and reporting duties (legal) [306]. Google's threat intelligence group reported in April 2025 that operations were expanding in Europe and that, with employee-owned devices, trails such as "corporate laptop shipping addresses and endpoint software inventories" are "unavailable" [307]. The mapping to functions is the authors' own.

No standard set of governance indicators was retrieved, so the authors propose four: time from first signal to a joint human resources, security, compliance and legal review; the share of incident reports that record a probability, a confidence level and at least one alternative hypothesis; the share of reports meeting the 24-hour, 72-hour

and one-month windows; and the share of contractors and remote hires with independent identity verification (Tier 3).

8.7 Spain

Spain combines a political attribution chain with an uneven technical one. Table 26 sets out the instruments that bear on detection and attribution.

Table 26. Spanish instruments for detection, attribution and reporting

Instrument	Date	What it provides	Gap in the text read
Disinformation procedure (Orden PCM/1030/2020)	Approved 6 October 2020; BOE 5 November 2020	Four levels; the National Security Council handles public attribution to a third State	Names no detection sources or attribution method; omits the CCN
Foro against disinformation campaigns (Orden PCM/541/2022)	Constituted 27 October 2022	Advisory body attached to the DSN; nine initiatives and over 70 experts in 2025	Outputs only; no outcome metrics found
National strategy procedure (Orden PJC/248/2025)	Approved 28 January 2025; BOE 15 March 2025	Drafting led by the Comisión Permanente; autonomous communities consulted	Sets no deadlines
CCN and CCN-CERT (Real Decreto 421/2004)	Functions page dated 6 June 2018	Guides, training, certification; response to threats to public-sector systems	Page does not mention attribution or threat reporting
NIS2 transposition bill	Presented 14 January 2025; Court referral 8 July 2026	A National Cybersecurity Center as crisis-management authority	Not yet law; Commission seeks a lump sum and daily penalties

Source: Authors' compilation from [308, 309, 310, 311, 312, 313].

The 2020 procedure puts monitoring and early warning at Level 1, with the intelligence service among the participants, and names the Secretaría de Estado de Comunicación as the single point of contact with the EU Rapid Alert System [308]. No public CNI product on attribution was found, so the service's role is known only from the procedure and the cybersecurity bill (Tier 3). That fits the pattern of Section 4: the decision to name a State sits at the top, and the technical assessment that feeds it is not described.

NIS2 leaves Spanish operators between EU law and enforceable national duties. The Commission sent a letter of formal notice on 28 November 2024 and a reasoned opinion on 7 May 2025, and on 8 July 2026 referred Spain, France, Ireland and the

Netherlands to the Court of Justice [313]. *Inference*. Article 23-style reporting, supplier scrutiny and board accountability are likely to bind Spanish operators once the law passes, so early alignment lowers the transition cost; this is a judgment about timing and no legal opinion (Tier 3). Spain's own record also supplies the case for the procedure above, the blackout of 28 April 2025 (Section 7). Annex D lists the Spanish notification channels now in force, which come from earlier law (D.2).

9. Political economy: why deniable operations are chosen, tolerated and rarely answered in kind

Sections 6 to 8 describe what operators do and how defenders assess it. This section asks why the pattern persists: why initiators choose deniable means, and why targets and their allies tolerate them and seldom answer in kind. The evidence base is thinner here than anywhere else in the manual. No study read measures incentives directly. The sources on motive are mostly abstracts or a publisher's description, and the only hard data are dated records of sanctions and delays. Every synthesis below is labeled *Inference* and carries Tier 3.

9.1 What the initiating actor gains

Six sources speak to the initiator's calculus, and each was read at limited depth (Table 27).

Table 27. What the sources say about the initiating actor's incentives

Source	Claim about incentives	Reading limit	Tier
Carson (2018)	Powerful states keep involvement secret to manage escalation and avoid domestic pressure	Publisher's description only	2
Cormac and Aldrich (2018)	Technology and contractors weaken plausible deniability; leaders accept "implausible deniability"	Abstract only	2
Eisenstadt (2021)	For Iran, the standoff that proxies provided mattered more than deniability	One analyst, one actor	2
Lindsay (2015)	Anonymity is easy for low-value attacks; attribution cost falls as target value rises	Argued for cyberattack; untested here	2
Kydd (2026)	A false flag is rational if cheap, with a credible culprit and a willing punisher	Abstract only; silent on relative strength	2
Levin (2025)	Low-credibility false flags raised third-party approval of war-seeking states and cut backing for sanctions	Abstract only, read on the author's page	2

Source: Authors' compilation from [8, 9, 10, 11, 63, 140].

The record adds the price actually paid in time. A first public attribution by a government came at most 2 days after the event in Crimea, 8 days after it in Salisbury, 2 years 8 months after it for Olympic Destroyer and about 10 years after it for TV5Monde, and a final finding or admission about seven weeks after it in Crimea and about 7 years 9 months after it in Salisbury (Section 6.5). No source shows an operator stopping because of an indictment or a sanction, and takedowns produced substitution (Section 5). Prigozhin acknowledged interference outright, through Concord's press service, on 7 November 2022. That is the "acknowledged" value of the attribution posture of Section 5, and it shows that concealment is not always the aim.

Inference. Read together, the sources describe deniability as inexpensive insurance and no binding constraint. The delay between act and answer is long, the answer falls on individuals and firms, and the accused denies. If Eisenstadt is right that standoff outweighed deniability for Iran, and Cormac and Aldrich that denials are believed less, the dividend lies in delay and doubt. The selection effect of Section 6 applies: only exposed operations are counted, so no source can say how many were chosen, abandoned or never found (Tier 3).

9.2 What attribution costs the target and its allies

Naming an actor is a political act that no rule requires (Section 4), and the costs fall on the namer. Dylan and Maguire (2022) list sources lost once exposed, narrowed

escalation options and credibility damage when claims lack support, and they warn that public use of intelligence can press officers toward "simpler, more alarming judgements" in a "perilous environment" shaped by the Iraq legacy [154]. A wrong call is expensive too. In the Iberian blackout speculation preceded evidence and the court, the government committee and ENTSO-E each reported no cyber cause (Section 7), and in Romania the dispute concerned the procedure and timing of acting on intelligence (Section 8).

Allies do not follow automatically. Levin (2025) reports that third-party publics grew more approving of war-seeking states after low-credibility false flags and less supportive of sanctions [63]. Carson (2018) argues that rival powers usually detect covert interventions but may stay silent because they share an interest in limiting the war [10]. The accused, for its part, can turn the label into an answer, as officials did with the pre-emptive "false flag" statements of Section 7, and Schiff, Schiff and Bueno (2025) found in five experiments with 15,287 US adults that politicians who falsely called a real scandal fake gained support [267].

Inference. The cost of naming is front-loaded and certain: credibility, sources and domestic politics. The benefit is delayed and diffuse, a deterrent effect that no source measured. A government can therefore rationally choose silence, a statement short of attribution, or a sanction on individuals (Tier 3).

9.3 The EU listing rounds as a case

The hybrid sanctions regime is the one institutional response with a dated record. Table 28 sets out the rounds located and the Council's printed totals.

Table 28. EU hybrid-threats sanctions: rounds located and printed totals, 2024 to 2026

Date	Action	Printed total	Reading limit
16 Dec 2024	First listings: 16 individuals, 3 entities	Not printed	Council release, read through a copy
20 May 2025	21 individuals, 6 entities; sectoral measures	Not printed	Press-release title only
15 Jul 2025	9 individuals, 6 entities	47 and 15	EEAS page
3 Oct 2025	One-year prolongation of the regime	Not printed	Press-release title only
15 Dec 2025	12 individuals, 2 entities	59 and 17	Council release read
29 Jan 2026	6 individuals, information manipulation	65 and 17	Copy of Council release
16 Mar 2026	4 individuals; Decision (CFSP) 2026/646	69 and 17	EEAS page
21 Apr 2026	2 entities, Euromore and Pravfond	69 and 19	Council release read
24 Sep 2026	1 individual; Decision (CFSP) 2026/2164, recorded as entry 81	Not printed	Decision text; Section 4
8 Oct 2026	Renewal to 9 October 2027; no new listings	80 and 20	EEAS page; press report

Source: Authors' compilation from [90, 104, 105, 292, 293, 314, 315, 316, 317, 318, 319]. Printed totals are individuals and entities as the sources state them.

Three features stand out. The first is pace: eight listing rounds were located in the 21 months from 16 December 2024 to 24 September 2026, most of them small. The second is arithmetic. Summing the rounds gives 46 individuals by 15 July 2025 against a printed 47. After 24 September 2026 the sum is 69, or 70 if the last printed total of 69 is advanced by the September listing, against 80 printed two weeks later, so 10 to 11 individuals and one entity sit in rounds the review did not find, or the totals rest on different bases. The 24 September listing is itself recorded as entry 81 beside a stated total of 80 individuals, and numbering against delistings is unresolved. This manual therefore gives 80 individuals and 20 entities as the Council's stated total at renewal and asserts no count of its own.

The third feature is the grounds. The December 2025 round names analysts and influencers "closely linked to or affiliated with the Kremlin's policy and messaging apparatus", members of GRU Unit 29155, and a Kaliningrad electronic warfare battalion linked to GPS failures in several Member States [292]. The September 2026 round lists Xenia Fedorova, a former head of RT France, for narratives said to be aligned with those of the Russian authorities; that account comes from a Danish legal-information summary of the Council act, so it is secondary [320].

Inference. Listing individuals and firms lets the Council answer without a finding of State responsibility, which keeps the political cost of naming a State out of the decision (Section 4). The record shows steady increments, which fits an instrument built for annual renewal. No source tests whether any listing changed an operator's conduct (Section 8), and the sources report no legal challenge or civil-liberties assessment of the speech-adjacent grounds; General Court challenges were not checked (Tier 3).

9.4 Who decides

Where attribution sits matters. The Council's overview of hybrid threats, as read, describes no formal attribution process, Spain reserves public attribution of a campaign to a third State for its National Security Council, and the United Kingdom and France publish unit-level or campaign-level attributions (Section 8). VIGINUM is the body reviewed that pairs published criteria with an independent ethical and scientific committee reporting to the Prime Minister. Csernaton and Wanless (2026) name threat inflation and the securitizing of civic initiatives as hazards, and observe that many European hybrid-threat offices "lack the capacities and authority required" [302].

Inference. A chain of command that detects, attributes and sanctions with no outside check has an incentive to widen its grounds, and an equal incentive to stay silent when a call is politically costly. The sources show the wide grounds (Table 28) and the multi-year silences (Section 6). None measures institutional capture directly, so it remains a hypothesis (Tier 3).

9.5 What could change the calculus

Three levers follow from the pattern, and the sources test none of them. The first raises the initiator's cost of ambiguity by shortening the interval between act and answer. Where an institution keeps a register, the latest step is days to three months old, and in the other rows of Table 9 the latest datum is from 1 month to 79 years old (Section 4). The second lowers the target's cost of naming by sharing technical evidence and vocabulary across allies while leaving the decision to name a State national, which every institution reviewed insists on. The third binds the attributing institution through published criteria and independent review, safeguards whose effect on false positives is untested (Section 8).

Assessment

The sources support a narrow claim: deniability pays in delay and doubt, and naming an actor is costly for the namer. They do not show that any measure has changed an operator's conduct. A mechanism that moved only the first lever would invite the over-attribution of Section 7. Section 10 sets out options for each lever with an explicit rating of political viability.

10. Policy options: the hammer and the carrot

Section 9 identified three levers: raise the initiator's cost of ambiguity, lower the target's cost of naming, and bind the institution that attributes. This section turns them into options. Each is a technical proposal assembled from instruments already in the record, and each carries a viability rating with the reason and the main obstacle (Table 29). High means the instrument is in force and was used in 2025 or 2026. Medium means it is adopted in part or needs one identifiable change. Low means no precedent was found or the obstacle is structural. The scale proxies political viability through precedent and does not forecast any Member State's vote. The ratings are the authors' judgment (Tier 3). Few sources rate feasibility at all: CEPA says that "not all deterrence measures will be easily adopted", and ECFR gives no rating [321, 322].

Table 29. Twelve options, their viability and the main obstacle

Option	Type	Viability	Reason	Main obstacle
Keep renewing and widening EU hybrid listings	Hammer	High	Regime renewed in October 2025 and again in October 2026	Speech-adjacent grounds invite challenge; no outcome evidence
Infringement procedures on late transposition of resilience and sanctions law	Hammer	High	Procedures on NIS2 and Directive 2024/1226 already run	Slow: 1 year 8 months from NIS2 deadline to Court referral
DSA incidents and crisis protocol for platforms	Hammer	Medium	Legal tools exist; the Democracy Shield plans the protocol	Output-only metrics; transatlantic frictions noted in the sources
Article 4 consultations first, Article 5 ambiguity kept	Hammer	Medium	Hybrid actions could trigger Article 5 since 2016; Article 4 requested 10 September 2025	Ankara 2026 text is silent on Article 4, so the signal is uncertain

Option	Type	Viability	Reason	Main obstacle
Secondary sanctions, seizures, offensive cyber, information operations	Hammer	Low	CEPA admits escalation risk; CSIS links the measures to peace talks	Unclear legal basis at sea and in cyberspace
Centrally managed resilience grants and technical assistance	Carrot	Medium	NATO's 1.5 percent line, NDICI and Democracy Shield programs adopted in principle	Competes with core defense spending; no effect metrics yet
Assisted-transition pacts tied to resilience milestones	Carrot	Low	ISDO proposal in another domain (June 2026 report); none found for hybrid resilience	Needs conditionality that Member States will accept
Time-limited problem cells and national ecosystem assessments	Carrot	Medium	Lower institutional cost than new agencies (Carnegie proposal)	Capacity gaps; threat inflation without written safeguards
Hybrid CoE branch in Spain (ECFR proposal)	Carrot	Low	No Spanish government position found	Needs a host government and Participating States; may add sprawl
Shared attribution vocabulary and evidence-package format	Enabler	Medium	ICD 203 and PHIA align by row and differ at band edges; the shared package is a proposal	Each State keeps the decision to name a State
Outcome indicators and dated milestones for EU programs and Spain's strategy	Enabler	Medium	ECA recommended indicators in 2021; Spain's procedure sets no deadlines	Not adopted since 2021: the 2021 findings still describe the 2026 reports
Pooled political attribution by several States	Enabler	Low	Technical convergence is feasible, but naming a State stays national	Every institution reviewed treats attribution as sovereign

Source: Authors' assessment (Tier 3) built on [34, 101, 104, 108, 109, 199, 281, 291, 294, 302, 310, 321, 322, 323, 324, 429].

10.1 The hammer: legal and financial options

The first two options use instruments that already run. The hybrid regime, set up on 8 October 2024, freezes assets, bars making funds available and imposes travel bans on individuals [71], and Table 28 shows its use. Infringement works on a slower clock. NIS2 reasoned opinions reached 19 Member States on 7 May 2025, and four were referred to the Court of Justice on 8 July 2026, with the Commission asking for a lump sum and daily penalties [112, 313]. The sanctions-crime directive follows the same rhythm

(Section 4). Both instruments rate High because they act on conduct and need no new law, and neither has an outcome evaluation (Section 8).

The harder options sit at the edge of what the sources support. CEPA urges pairing deterrence by denial with deterrence by punishment, lists asset freezes, Magnitsky-style frameworks, expulsions of intelligence-linked diplomats, prosecutions and coordinated takedowns, and concedes that attribution makes punishment "a challenging undertaking" and that any response "must be proportionate" [321]. CSIS adds secondary sanctions, seizures of shadow fleet vessels and offensive cyber operations, the most hawkish position retrieved, with no feasibility or legal analysis in the version read [199]. These options rate Low in the EU. Acting on a wrong attribution carries legal risk of its own, since countermeasures lie only against a State responsible for a wrongful act (Section 4).

Inference. Every punitive option should travel with the evidential thresholds of Section 8, because the cost of a wrong call falls on the State that made it (Tier 3).

10.2 The carrot: institutional incentives

The carrot exists mostly as spending lines and voluntary bodies. NATO's Hague text allows up to 1.5 percent of GDP a year for critical infrastructure, networks, civil preparedness and resilience, among other purposes, and NATO's counter-hybrid support teams assist Allies on request [291, 324]. The EU's Hybrid Rapid Response Teams support Member States, partners and CSDP missions, with no public deployment record located [290]. The Democracy Shield adds EUR 5 million for fact-checkers, media resilience programs and a Centre for Democratic Resilience with voluntary participation [323]. Spain shows what is missing: the strategy procedure sets no deadlines, the Foro reports outputs and no outcomes, and the NIS2 bill is not law (Section 8).

ISDO's June 2026 report on discrimination in Europe proposes a template for pairing the two sides. It keeps legal enforcement in force and adds positive conditionality, with funds redirected under supervision, and assisted governance transition pacts in which a specialist agency temporarily co-manages overwhelmed operations. Both additions are proposals of that report (its Section 11.3) [429]. An analogue here would pair listings and legal consequences for perpetrators with centrally managed technical assistance and resilience grants for exposed Member States and operators. No source read proposes exactly this pairing, and none offers evidence on conditionality as a lever in this domain, so it is an ISDO proposal and no cited result

(Tier 3). The southern Hybrid CoE that ECFR proposes for Spain, focused on Spanish-language information ecosystems, would need a host government and Participating States [322]. Carnegie Europe prefers time-limited cells, for example for an electoral period or a Baltic infrastructure cell, to a permanent "supersilo" [302].

10.3 Enablers, sequence and a caution

Three options concern the quality of attribution and the accountability of those who attribute. A shared vocabulary is cheap: the two official scales of Table 21 align row by row, and naming the scale costs nothing. Pooled political attribution is a different matter, because every institution reviewed treats the decision to name a State as sovereign (Section 4), so only the technical layer can be shared. Outcome indicators are the cheapest option, and no source read opposes them. The Court of Auditors recommended in 2021 a monitoring and evaluation framework with measurable indicators, regular reporting and funding transparency, and its findings still describe the 2026 reports (Section 8) [294]. Both enablers rate Medium: the vocabulary option includes an evidence-package format that does not yet exist, and the indicators have waited since 2021.

Assessment

Every option in Table 29 can be delivered as an output, and none has an outcome metric yet. That makes evaluation the cheapest first step and a candidate precondition for the rest. Options that raise the cost of ambiguity without the thresholds of Section 8 would reproduce the over-attribution of Section 7, and options that only lower the cost of naming would leave the initiator's calculus unchanged.

Sequence follows from the ratings and from what each option enables. Listings and infringement, the two High options, already run and continue. Among the options still to be adopted, evaluation and a shared vocabulary come first: each needs one identifiable change, each costs little, and each enables the rest. Resilience grants follow, and assisted-transition pacts wait for a pilot. The Low options stay on a watch list until their obstacles change. These are technical proposals for an analytical audience, and each can be dropped without affecting the others.

11. Conclusions and research agenda

Sections 2 to 10 followed one question from definition to policy: what can public evidence say about operations built to be denied? The answer has five findings, a short list of limits and five questions that no source read could settle.

11.1 Findings

The strict false flag is rare in the documented record. Under the evidence rule of Section 3, the documented cases are historical: Mukden, Gleiwitz and, through secondary sources only, Mainila and the Lavon Affair. Olympic Destroyer and the 2020 Proud Boys emails are the contemporary candidates, and both stay contested. Most of the contemporary operations reviewed (the graded register runs from 2018 to 2026) rely on deniability by distance, through proxies, contractors and personas, and place blame on no third party (Section 6). Attribution comes in layers. A first public attribution by a government came at most 2 days after the event in Crimea and 8 days after it in Salisbury, but 2 years 8 months after it for Olympic Destroyer and about 10 years for TV5Monde, and a final finding or admission took about seven weeks in Crimea and about 7 years 9 months in Salisbury (Section 6.5).

The law reaches the elements of a deniable operation and misses its design. Each regime applies above a threshold that such an operation is built to stay under, attribution is a sovereign political act, and the latest ruling located on enemy colors dates from 1947 (Section 4). EU person-based sanctions bypass the State-responsibility threshold, and they are the one response with a dated record of rounds (Section 9).

Access, more than rhetoric, separates settled claims from open ones. Of the fifteen claims selected in Section 7, a set that is no sample, eight are Refuted, each after an inquiry, court or inspection body examined the specific claim, and seven stay Alleged. Where access was denied or the mandate excluded attribution, the accusations stayed open.

Defenders have standards and little proof that they work. ICD 203 and the UK yardstick demand alternative hypotheses and calibrated language, yet the one randomized test of ACH on practicing analysts that was located found no gain in accuracy, and the sources read show no outcome evaluation of any institutional countermeasure (Section 8). The two-sided procedure of that section is the authors' proposal and is untested, and so is Annex D, which turns it into protocols.

Finally, deniability pays in delay and doubt, and naming an actor costs the namer (Section 9). Section 10 rates twelve options, two of them High, and puts evaluation first among those still to be adopted: every option can be delivered as an output, and none has an outcome metric.

11.2 Limits

The record is a sample of operations that were detected and, with few exceptions, publicly attributed, so no rate of staged blame can be drawn from it. The fifteen claims of Section 7 are likewise a selection, and eight Refuted of fifteen is no rate of false accusations. Nearly every attribution is one authority's account, many sources were read at abstract level or through a single route, and the grades are the authors' judgment. The thresholds of Table 23 have not been tested on cases of known outcome (Section 8.3). Three counts could not be reconciled and are asserted nowhere: the EU hybrid-listing totals (Section 9), the States Parties to the Kampala amendments and the membership of the Special Tribunal (Section 4). Every status ends on October 11, 2026.

11.3 Open questions

The detection research points to five priorities (Table 30).

Table 30. Five priority research questions and what the record holds now

No.	Open question	What the record holds now
1	What false-positive and false-negative rates attach to metadata, geolocation and signature-mismatch indicators?	Single detection cases; no error rate for any class (Section 8)
2	Does public attribution or sanction change hostile behavior, separating detection effects from activity effects?	Output counts; no operator shown stopping after an indictment (Sections 5 and 8)
3	Do structured techniques improve attribution accuracy under realistic deception hypotheses?	The one randomized test located: 50 analysts; no accuracy gain, lower consistency (Section 8)
4	Does pre-bunking build discernment or generic skepticism, over months and not minutes?	Short-term experiments; no evidence on durability (Section 8)
5	How can legal proof standards and political attribution be reconciled for seabed and infrastructure cases?	Eagle S: jurisdiction affirmed in 2026; no State link found (Section 4)

Source: Authors' agenda from the detection research and Sections 4, 5 and 8 [91, 284, 298].

11.4 What future editions should add

A later edition should re-read the primary documents behind each Tier 2 attribution and add the accused parties' replies that this edition could not retrieve. It should refresh every status in Annex B.1: the Nord Stream trial opened on 14 October 2026, three days after this edition's evidence cutoff (Section 6.5), and the Eagle S appeal deadline is 26 October 2026. It should test the protocols of Annex D in tabletop exercises and on closed incidents, publish the call-and-outcome register that annex proposes, and refresh the Spanish notification channels once Spain transposes NIS2. It should also review Kilcullen and Rupert Smith and, where sources allow, disaggregate the exposure and belief evidence of Section 7 by gender and age. Until then the working rule of Section 8 stands: state an evidential threshold in both directions, and keep to it.

Closing note

This second edition is dated October 16, 2026, and its evidence cutoff is October 11, 2026. Compared with the edition of October 12, 2026, it adds Annex D, which holds the forensic protocols, the verification guide, the three response flows and the record templates. It corrects a row and a sentence on ICD 203 in Section 8 and the source notes of Tables 21, 25 and 52, restates the Berkeley Protocol and NIST passages of Section 8 after reading those texts directly, and adds fourteen glossary terms. Short cross-references to Annex D now appear in the Abstract, the Executive summary and Sections 1, 8 and 11. Bibliography entries [281], [285], [303] and [304] are updated, and entries [403] to [416] are new. Before the corrections described next, the rest of Sections 1 to 11 stood as first published.

Corrections of October 22, 2026. After an internal critical review, the authors corrected this edition. It keeps its edition date of October 16, 2026 and its evidence cutoff of October 11, 2026. The changes that alter a count, a grade, a rating or a reading are these.

- **Counts and grades.** A test is met only when it is graded Y. On that rule 20 of the 38 incidents meet no test, six meet at least two and one meets all three, and the count that also credits a partly met test, 17 meeting all three, stands beside it. Grade U, which had covered both evidence that is not public and sources this review did not retrieve, is split in two, and six cells change from U to R (Annex C.1.2). The other 108 grades of the three tests are unchanged.

- **Claims.** Row N13 of the claims register is split into N13a and N13b, both graded Alleged, so the register holds fifteen claims, eight Refuted and seven Alleged, where it held fourteen. The rule that separates Refuted from Alleged is written down, and the Key Bridge claim now cites the FBI Baltimore statement [432].
- **Cases.** The list of documented and contested false flags adds the Lavon Affair, known through secondary sources and with a disputed chain of command, and names the 2020 Proud Boys emails as a contested contemporary candidate beside Olympic Destroyer. The notes to Tables 16 and 36 state the source basis of Mainila, Lavon and TPAJAX.
- **Law and practice.** Table 9 gives exact dates for each obligation and adds a second clock, the age of the latest datum at the cutoff, and Figure 2 now charts that age, which does not depend on the age of the norm. The interval from event to authoritative statement is replaced by two intervals counted from the event, one to the first public attribution by a government and one to the final finding or admission (Sections 4.7, 6.5, 7.4, 8.3, 9.1 and 11.1, Annex D.3.5 and Table 1). The new Figure 4 plots them.
- **Policy ratings.** Two of the twelve options in Table 29 move from High to Medium, because they do not meet the scale's own definition of High, an instrument in force and used in 2025 or 2026. The count is now two High, six Medium and four Low. Section 10 states that the scale proxies political viability through precedent, the sequence of Section 10.3 is derived again from the new ratings, and the June 2026 ISDO report that proposes one of the options is cited as a proposal [429].
- **Court register.** Docket numbers, courts and statuses in Tables 44 and 46 were completed or corrected from the sources added as entries [417] to [428], consulted on October 22, 2026. A number that could not be confirmed is marked as not located or unconfirmed, and none is printed in its place.
- **Limits and disclosure.** Sections 1.2 to 1.4 and 11.2 and Annex C now state where the evidence stops. Absence claims hold for the sources read, the incident register and the claims set are selections and no samples, and the 114 grades are one reader's judgment, to be coded blind by both authors, so the counts may change. The commissioning list of 34 sources is described in neutral terms, with the audit of that list moved to Section 1.4. Annex C.1.4 discloses that the AI model used in drafting comes from the vendor behind incident I29, and that an outside reader has not yet checked that incident.

- **Smaller corrections.** Arithmetic, labels, captions and source notes are corrected, for example in Sections 2.1, 4.5, 5.7, 6.1, 8.1 and 9.3, and the EU listings after Salisbury, on January 21, 2019 [230], are timed at 10 months 17 days after the event.
- **Additions.** Figure 4 (intervals from the event to attribution and to a final finding), Figure 5 (the procedure of Section 8.3 as a flow) and Annex E (a Spanish executive summary and a bilingual glossary of the 58 terms) are new. Bibliography entries [417] to [437] are new, and entries [63], [92], [93], [101], [166], [196] and [221] are updated. The open data list every edit.

Evidence gathered for the manual still stops at October 11, 2026, with the exceptions that the text marks: court statuses and register balances dated after that day (Section 6.5 and the notes to Tables 44 and 46), and the readings of October 22, 2026 that the bibliography records.

The authors thank the ISDO Executive Board for commissioning and supporting the work, and the public institutions and open-data communities whose published records, court documents and datasets made verification possible. Any error that remains is the authors' own. The manual is published by ISDO, Zaragoza, under the CC BY-NC 4.0 license.

Recommended citation: Rodríguez, Á., & Sainz, S. (2026). Manual on False-Flag Operations and Hybrid Warfare in the Context of Asymmetric Warfare (2nd ed., corrected October 22, 2026). ISDO Security and Intelligence Department. October 16, 2026. <https://isdo.ch>

Bibliography

Reading note. Entries carry markers where the research notes were incomplete. "Title as shown in the URL" means the title was taken from the web address and not checked against the page. "Title not recorded" and "date not recorded" mean the notes hold no value, and the entry then gives a descriptive label. "n.d." means no date was found on the page, and "year not shown" means the source gives none. "Read through a tool summary" means a summarizing tool returned the page text. The markers describe the entry's data and say nothing about the quality of the source.

- [1] Mack, A. (1975). Why Big Nations Lose Small Wars: The Politics of Asymmetric Conflict. *World Politics* 27(2), 175-200. DOI 10.2307/2009880. <https://www.cambridge.org/core/journals/world-politics/article/abs/why-big-nations-lose-small-wars-the-politics-of-asymmetric-conflict/90583542E0F98B15B0A2C37D390C9C41>
- [2] Arreguín-Toft (2005). How the Weak Win Wars: A Theory of Asymmetric Conflict. Cambridge University Press; read through the Belfer Center summary. <https://www.belfercenter.org/publication/how-weak-win-wars-theory-asymmetric-conflict>
- [3] Münkler (2003, March). Article on asymmetric war (title not recorded). *International Review of the Red Cross* 85(849). https://www.icrc.org/en/doc/assets/files/other/irrc_849_munkler.pdf
- [4] Kaldor (2013). In Defence of New Wars. *Stability* 2(1). <https://www.stabilityjournal.org/articles/10.5334/sta.at>
- [5] Mao Zedong (1938). On Protracted War. Selected Works, vol. 2, Marxists Internet Archive. https://www.marxists.org/reference/archive/mao/selected-works/volume-2/mswv2_09.htm
- [6] Qiao Liang and Wang Xiangsui (1999, February). Unrestricted Warfare. PLA Literature and Arts Publishing House; FBIS translation. <https://www.c4i.org/unrestricted.pdf>
- [7] Metz and Johnson (2001, January). Monograph on asymmetry (full title not recorded). Strategic Studies Institute, U.S. Army War College Press. <https://press.armywarcollege.edu/monographs/121/>
- [8] Cormac, R., and Aldrich, R. J. (2018). Grey is the new black: covert action and implausible deniability. *International Affairs* 94(3), 477-494. DOI 10.1093/ia/iiy067. <https://nottingham-repository.worktribe.com/output/961763/grey-is-the-new-black-covert-action-and-implausible-deniability>
- [9] Eisenstadt (2021). Iran's gray zone strategy: cornerstone of its asymmetric way of war (title as shown in the URL). *PRISM* 9(2), NDU Press. <https://ndupress.ndu.edu/Media/News/News-Article-View/Article/2541911/irans-gray-zone-strategy-cornerstone-of-its-asymmetric-way-of-war/>
- [10] Carson (2018, September 25). Secret Wars: Covert Conflict in International Politics. Princeton University Press; publisher description only. <https://press.princeton.edu/node/6815>
- [11] Kydd, A. H. (2026, April 13). False flags: a strategic approach. *Quarterly Journal of Political Science* 21(2), 129-151. DOI 10.1108/QJPS-09-2024-0122; abstract read through the OpenAlex record. <https://doi.org/10.1108/QJPS-09-2024-0122>
- [12] Hoffman (2007). Monograph on hybrid wars (title not recorded). Potomac Institute for Policy Studies. https://www.potomacinstitute.org/images/stories/publications/potomac_hybridwar_0108.pdf

- [13] Mazarr (2015). Monograph on gray-zone strategy (title not recorded). U.S. Army War College Press. <https://press.armywarcollege.edu/monographs/428/>
- [14] Morris et al. (2019, June 27). RAND Research Report RR2942 on gray-zone coercion (full title not recorded). RAND Corporation. https://www.rand.org/pubs/research_reports/RR2942.html
- [15] Cullen and Reichborn-Kjennerud (2017, January). Countering Hybrid Warfare Project: Understanding Hybrid Warfare. Multinational Capability Development Campaign; NUPI. https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/647776/dar_mcdc_hybrid_warfare.pdf
- [16] Kofman and Rojansky (2015, April). Kennan Cable No. 7 (title not recorded). Kennan Institute, Wilson Center. <https://www.wilsoncenter.org/sites/default/files/media/documents/publication/7-KENNAN%20CABLE-ROJANSKY%20KOFMAN.pdf>
- [17] Stoker and Whiteside (2020). Blurred Lines: Gray-Zone Conflict and Hybrid War. Naval War College Review 73(1), Article 4; abstract read. <https://digital-commons.usnwc.edu/nwc-review/vol73/iss1/4/>
- [18] NATO Review (2015, May 7). Hybrid war, does it even exist? NATO; authorship (Damien Van Puyvelde) per the University of Glasgow repository record, <https://eprints.gla.ac.uk/147728>. <https://www.nato.int/docu/review/articles/2015/05/07/hybrid-war-does-it-even-exist/index.html>
- [19] Renz (2016). Russia and 'hybrid warfare'. Contemporary Politics 22(3), 283-300. DOI 10.1080/13569775.2016.1201316; abstract read. <https://nottingham-repository.worktribe.com/output/793609>
- [20] Renz and Smith (2016). Russia and Hybrid Warfare: Going Beyond the Label. Aleksanteri Papers 1/2016, Aleksanteri Institute, University of Helsinki. <https://helda.helsinki.fi/server/api/core/bitstreams/9514b166-0249-42a4-a408-9195e7d32292/content>
- [21] Galeotti (2018, March 5). I'm Sorry for Creating the 'Gerasimov Doctrine'. Foreign Policy. <https://foreignpolicy.com/2018/03/05/im-sorry-for-creating-the-gerasimov-doctrine/>
- [22] Person, R., Kulalic, I., and Mayle, J. (2024, July). Back to the Future: The Persistent Problems of Hybrid War. International Affairs 100(4), 1749-1761. DOI 10.1093/ia/iaae131; abstract read. <https://doi.org/10.1093/ia/iaae131>
- [23] Radin (2017, February 23). Hybrid Warfare in the Baltics: Threats and Potential Responses. RAND Research Report RR-1577-AF. https://www.rand.org/pubs/research_reports/RR1577.html
- [24] Gerasimov (2013, February 27). Article in VPK Kurier (title not recorded); English translation by Coalson, Military Review, January-February 2016, pp. 23-29. https://www.armyupress.army.mil/Portals/7/military-review/Archives/English/MilitaryReview_20160228_art008.pdf
- [25] Russian Federation (2014). Military Doctrine of the Russian Federation. Kremlin-hosted PDF; read through a summarizing tool. <http://static.kremlin.ru/media/events/files/41d527556bec8deb3530.pdf>
- [26] Russian Federation (2021, July 2). Strategy of National Security of the Russian Federation (Decree No. 400); unofficial English translation hosted by the University of Surrey. https://surrey.ac.uk/sites/default/files/2024-09/2021_Russia.pdf
- [27] Kania (2016, August). The PLA's latest strategic thinking on the Three Warfares (title as shown in the URL). China Brief, Jamestown Foundation. <https://jamestown.org/program/the-plas-latest-strategic-thinking-on-the-three-warfares/>

- [28] Washington Institute for Near East Policy (2020, January 7). Policy Focus 162: Operating in the gray zone, countering Iran's asymmetric way of war (title as shown in the URL). <https://www.washingtoninstitute.org/policy-analysis/operating-gray-zone-countering-irans-asymmetric-way-war>
- [29] NATO (2016). Warsaw Summit Communiqué. https://www.nato.int/cps/en/natohq/official_texts_133169.htm
- [30] NATO (2021). Brussels Summit Communiqué. https://www.nato.int/cps/en/natohq/news_185000.htm
- [31] NATO (2022). Strategic Concept 2022. <https://www.nato.int/content/dam/nato/webready/documents/publications-and-reports/strategic-concepts/2022/290622-strategic-concept.pdf>
- [32] European Commission and High Representative (2016, April 6). Joint Framework on countering hybrid threats, JOIN(2016) 18 final. EUR-Lex. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=JOIN:2016:18:FIN>
- [33] Council of the European Union (2022, March). Strategic Compass. EEAS. https://www.eeas.europa.eu/sites/default/files/documents/strategic_compass_en3_web.pdf
- [34] Council of the European Union (2026, March 16). Council Conclusions on advancing the EU's capacity to counter hybrid threats, document 7349/26. <https://data.consilium.europa.eu/doc/document/ST-7349-2026-INIT/en/pdf>
- [35] U.S. Joint Staff (2018, March 16). Joint Concept for Integrated Campaigning; read through a summarizing tool. https://www.jcs.mil/Portals/36/Documents/Doctrine/concepts/joint_concept_integrated_campaign.pdf
- [36] U.S. Department of War (2026, January 23). 2026 National Defense Strategy; read as the USNI News excerpt of 24 January 2026. <https://news.usni.org/2026/01/24/2026-u-s-national-defense-strategy>
- [37] United Kingdom Government (2025, June 24). National Security Strategy 2025. https://assets.publishing.service.gov.uk/media/685ab0da72588f418862075c/E03360428_National_Security_Strategy_Accessible.pdf
- [38] Gobierno de España (2021, December 28). Estrategia de Seguridad Nacional 2021, Real Decreto 1150/2021. Boletín Oficial del Estado no. 314, 31 December 2021, BOE-A-2021-21884. <https://www.boe.es/boe/dias/2021/12/31/pdfs/BOE-A-2021-21884.pdf>
- [39] Ministerio de Defensa de España (2020, August 4). Directiva de Política de Defensa 2020. Boletín Oficial del Ministerio de Defensa no. 159, 6 August 2020; read through a summarizing tool. <https://working.mde.es/Galerias/defensadocs/directiva-politica-Defensa-2020.pdf>
- [40] Pynnöniemi (2021, May). The concept of hybrid war in Russia. Hybrid CoE Strategic Analysis 27. https://www.hybridcoe.fi/wp-content/uploads/2021/05/20210518_Hybrid_CoE_Strategic_Analysis_27_The_concept_of_hybrid_war_in_Russia_WEB.pdf
- [41] NATO (2024). Washington Summit Declaration. https://www.nato.int/cps/en/natohq/official_texts_227678.htm
- [42] National Security Council (1948, June 18). NSC 10/2, covert operations directive. Foreign Relations of the United States 1945-1950, Intelligence, doc. 292. <https://history.state.gov/historicaldocuments/frus1945-50Intel/d292>
- [43] United States Code (n.d.). 50 U.S.C. 3093, covert action provisions. Cornell Legal Information Institute. <https://www.law.cornell.edu/uscode/text/50/3093>

- [44] EEAS (2025, March 19). 3rd Report on FIMI Threats. European External Action Service. <https://www.eeas.europa.eu/sites/default/files/documents/2025/EEAS-3rd-ThreatReport-March-2025-05-Digital-HD.pdf>
- [45] Merriam-Webster (n.d.). False flag. Merriam-Webster Dictionary. <https://www.merriam-webster.com/dictionary/false%20flag>
- [46] San Remo Manual on International Law Applicable to Armed Conflicts at Sea (1994, June 12), paras 110-111; reproduction by the ICRC. <https://www.icrc.org/eng/resources/documents/article/other/57jmsu.htm>
- [47] NATO (2014, September 2). AJP-3.10.1, psychological operations doctrine, Edition B Version 1; withdrawn edition (UK-hosted copy). https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/450521/20150223-AJP_3_10_1_PSYOPS_with_UK_Green_pages.pdf
- [48] NATO (2020, March 6). AJP-3.10.2, Operations Security and Deception, Edition A Version 2; the copy read carries a withdrawal marking (UK-hosted). https://assets.publishing.service.gov.uk/media/69ccfbd59b2e6e135502d04a/ARCHIVE-AJP_3_10_2_Ops_and_Deception_EdA-O.pdf
- [49] U.S. Joint Chiefs of Staff (2006, July 13). Joint Publication 3-13.4, Military Deception; later editions not checked. https://edocs.nps.edu/dodpubs/topic/jointpubs/JP3/JP3_13_4_060713.pdf
- [50] Theohary, C. A. (2024, November 29). Defense Primer: Operations in the Information Environment, IF10771. Congressional Research Service; read on a mirror. https://www.everycrsreport.com/files/2024-11-29_IF10771_ce460ed092b161b6ddde2add47dd3ca25a3b479d.html
- [51] Rid, T., and Buchanan, B. (2015). Attributing Cyber Attacks. Journal of Strategic Studies 39(1), 4-37. DOI 10.1080/01402390.2014.977382; quotations taken from the authors' workshop version. https://www.sipa.columbia.edu/sites/default/files/2022-11/Cyber_Workshop_Attributing%20cyber%20attacks.pdf
- [52] Rid (2017, March 30). Written testimony to the Senate Select Committee on Intelligence. <https://www.intelligence.senate.gov/sites/default/files/documents/os-trid-033017.pdf>
- [53] Tallinn Manual 2.0 on the International Law Applicable to Cyber Operations (2017), table of contents (Rules 122-127); with the CCDCOE Cyber Law Toolkit page on perfidy and ruses of war. <https://www.obalkyknih.cz/file/toc/274573/pdf>
- [54] National Security Archive (2001, April 30). Operation Northwoods documents, including the Joint Chiefs memorandum "Justification for U.S. Military Intervention in Cuba" of 13 March 1962. <https://nsarchive2.gwu.edu/news/20010430/>
- [55] National Security Archive (2005, December 1). Gulf of Tonkin documents and the NSA historian Hanyok's study (EBB 132). <https://nsarchive2.gwu.edu/NSAEBB/NSAEBB132/press20051201.htm>
- [56] UK Government (2020, October 19). UK exposes series of Russian cyber attacks against Olympic and Paralympic Games (press release; FCDO and NCSC). <https://www.gov.uk/government/news/uk-exposes-series-of-russian-cyber-attacks-against-olympic-and-paralympic-games>
- [57] Additional Protocol I to the Geneva Conventions (1977), Arts 37-39. Text hosted by OHCHR. <https://www.ohchr.org/en/instruments-mechanisms/instruments/protocol-additional-geneva-conventions-12-august-1949-and>
- [58] ICRC (2005). Customary International Humanitarian Law, Volume I: Rules (Rules 57-65). <https://www.icrc.org/en/doc/assets/files/other/customary-international-humanitarian-law-i-icrc-eng.pdf>

- [59] Trial of Skorzeny and Others. General Military Government Court of the U.S. Zone of Germany (Dachau), 18 August to 9 September 1947; no docket number shown; final acquittal. Law Reports of Trials of War Criminals vol. IX (1949), pp. 90-93; ICRC Casebook. <https://casebook.icrc.org/case-study/united-states-military-court-germany-trial-skorzeny-and-others>
- [60] Kaspersky (2018, March). The devil's in the Rich header (Olympic Destroyer analysis; title as shown in the URL). Securelist. <https://securelist.com/the-devils-in-the-rich-header/84348/>
- [61] U.S. Department of Justice (2020, October 19). Six Russian GRU officers charged in connection with worldwide deployment of destructive malware (press release; indictment returned 15 October 2020, W.D. Pennsylvania). <https://www.justice.gov/opa/pr/six-russian-gru-officers-charged-connection-worldwide-deployment-destructive-malware-and>
- [62] NSA and UK NCSC (2019, October 21). Joint advisory on Turla group activity (U/OO/200914-19, PP-19-1354). https://media.defense.gov/2019/Oct/18/2002197242/-1/-1/0/NSA_CSA_TURLA_20191021%20VER%203%20-%20COPY.PDF
- [63] Levin, D. H. (2025, April 10 online). Article on low-credibility false-flag incidents (the commissioning list gives the title only as "How Low Can You Go? ... Low Credibility False Flag Incidents ..."; full title not recorded). International Interactions 51(3), 425-454. DOI 10.1080/03050629.2025.2489366; abstract read on the author's page. <https://www.dovhlevin.com/research-1>
- [64] Hague Convention IV (1907, October 18). Regulations respecting the Laws and Customs of War on Land, Arts 22, 23 and 24. Avalon Project, Yale Law School. https://avalon.law.yale.edu/20th_century/hague04.asp
- [65] United Nations (1945). Charter of the United Nations, Arts 2, 39 and 51 (chapters I and VII). <https://www.un.org/en/about-us/un-charter/chapter-1>
- [66] International Law Commission (2001). Articles on Responsibility of States for Internationally Wrongful Acts, Arts 4, 5, 8, 11, 25, 49-54. United Nations. https://legal.un.org/ilc/texts/instruments/english/draft_articles/9_6_2001.pdf
- [67] UN Group of Governmental Experts (2015, July 22). Report A/70/174, paras 13 and 28; UNIDIR extract. <https://unidir.org/files/2019-10/GGE-Recommendations-International-Law.pdf>
- [68] Council of the European Union (2017, June 7). Draft Council conclusions on the Cyber Diplomacy Toolbox, doc. 9916/17, paras 4-5 (full title not recorded); the adopted text of 20 November 2017 was not read. <https://data.consilium.europa.eu/doc/document/ST-9916-2017-INIT/en/pdf>
- [69] Council of the European Union (2024, November). Declaration on a Common Understanding of International Law in Cyberspace, doc. 15833/24. <https://data.consilium.europa.eu/doc/document/ST-15833-2024-INIT/en/pdf>
- [70] Council of the European Union (2019). Decision (CFSP) 2019/797, restrictive measures against cyber-attacks (cyber sanctions regime), Arts 1 and 4. EUR-Lex. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32019D0797>
- [71] Council of the European Union (2024, October 8). Decision (CFSP) 2024/2643, hybrid-threats sanctions regime (full title not recorded), Arts 1, 2 and 10. EUR-Lex. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32024D2643>
- [72] Rome Statute of the International Criminal Court (1998), Arts 8(2)(b)(vii), (xi), (e)(ix) and 8 bis; primary text not retrieved in the research. <https://www.icc-cpi.int/sites/default/files/RS-Eng.pdf>
- [73] Cyber Law Toolkit (n.d.). Perfidy and ruses of war. CCDCOE. https://cyberlaw.ccdcoe.org/wiki/Perfidy_and_ruses_of_war

- [74] Price (2020, February 21). Perfidy in cyberspace: the requirement for human confidence (title as shown in the URL). Harvard National Security Journal Online; not confirmed as peer reviewed. <http://journals.law.harvard.edu/nsj/2020/02/perfidy-in-cyberspace-the-requirement-for-human-confidence/>
- [75] International Law Commission (2001). Commentaries to the Articles on Responsibility of States for Internationally Wrongful Acts, Art 8 commentary paras 3-4; the extraction ended inside para 4. United Nations. https://legal.un.org/ilc/texts/instruments/english/commentaries/9_6_2001.pdf
- [76] International Court of Justice (1986, June 27). Nicaragua v. United States of America, merits, para 115; as presented in the ICRC Casebook. <https://casebook.icrc.org/case-study/icj-nicaragua-v-united-states>
- [77] International Court of Justice (n.d.). Military and Paramilitary Activities in and against Nicaragua (Nicaragua v. United States of America), case page, ICJ no. 70. <https://www.icj-cij.org/case/70>
- [78] International Court of Justice (n.d.). Bosnia and Herzegovina v. Serbia and Montenegro (genocide), case page, ICJ no. 91; judgment of 26 February 2007. <https://www.icj-cij.org/case/91>
- [79] Cassese, A. (2007). Article on the ICJ's treatment of the overall control test in the Bosnia genocide judgment (title not recorded). European Journal of International Law 18(4), 649-668. <https://academic.oup.com/ejil/article/18/4/649/453762>
- [80] European Court of Human Rights, Grand Chamber (2025, July 9). Ukraine and the Netherlands v. Russia, nos. 8019/16, 43800/14, 28525/20 and 11055/22; press release ECHR 173 (2025), Perugia copy. <https://diritti-cedu.unipg.it/wp-content/uploads/2025/07/GC-Corte-EDU-Ucrainan-e-Olanda-vs.Russia-press-release-9-luglio-2025.pdf>
- [81] International Court of Justice (n.d.). Armed Activities on the Territory of the Congo (Democratic Republic of the Congo v. Uganda), case page, ICJ no. 116; merits 19 December 2005, reparations 9 February 2022. <https://www.icj-cij.org/case/116>
- [82] Cyber Law Toolkit (n.d.). Countermeasures. CCDCOE. <https://cyberlaw.ccdcoe.org/wiki/Countermeasures>
- [83] Assembly of States Parties, International Criminal Court (2025, July 10). Press release PR-20250710 on the special session on the crime of aggression (title not recorded). <https://asp.icc-cpi.int/press-releases/PR-20250710>
- [84] Crime of Aggression website (page current to 2025, September 7). Status of ratification and implementation. <https://crimeofaggression.info/the-role-of-states/status-of-ratification-and-implementation/>
- [85] Council of Europe (2026). Special Tribunal for the Crime of Aggression against Ukraine: frequently asked questions. <https://www.coe.int/en/web/portal/special-tribunal-for-the-crime-of-aggression-against-ukraine-frequently-asked-questions>
- [86] International Centre for Defence and Security (2026, June 18). The Special Tribunal for Ukraine: turning the wheels of international justice (title as shown in the URL). <https://icds.ee/en/the-special-tribunal-for-ukraine-turning-the-wheels-of-international-justice/>
- [87] Cyber Law Toolkit (n.d.). Prohibition of intervention (revision oldid=2717). CCDCOE. https://cyberlaw.ccdcoe.org/wiki/Prohibition_of_intervention?oldid=2717
- [88] United Kingdom (2021, June 3). Application of international law to States' conduct in cyberspace: UK statement. GOV.UK. <https://www.gov.uk/government/publications/application-of-international-law-to-states-conduct-in-cyberspace-uk-statement/application-of-international-law-to-states-conduct-in-cyberspace-uk-statement>

- [89] Federal Government of Germany (2021, March). On the application of international law in cyberspace: position paper. Federal Foreign Office. <https://www.auswaertiges-amt.de/blob/2446304/32e7b2498e10b74fb17204c54665bdf0/on-the-application-of-international-law-in-cyberspace-data.pdf>
- [90] Council of the European Union (2024, December 16). Press release on the first listings under the hybrid-threats regime; GlobalSecurity copy. <https://www.globalsecurity.org/intell/library/news/2024/intell-241216-ec01.htm>
- [91] The Record (reported 2026, August 27). Finland appeals court revives case against Eagle S (title as shown in the URL). <https://therecord.media/finland-appeals-court-revives-case-against-eagle-s>
- [92] United Nations (1966, December 19). International Covenant on Civil and Political Rights, Arts 19 and 20(1); United Nations Treaty Series vol. 999, no. I-14668 (text read through a tool summary on October 22, 2026). <https://treaties.un.org/doc/publication/unts/volume%20999/volume-999-i-14668-english.pdf>
- [93] United Nations General Assembly (1970, October 24). Declaration on Principles of International Law concerning Friendly Relations and Cooperation among States in accordance with the Charter of the United Nations, A/RES/2625(XXV) (catalog record checked on October 22, 2026; text read through a secondary source). <https://digitallibrary.un.org/record/202170?ln=en>
- [94] Schmitt, M. (2021, October 3). Q&A about the status of Tallinn 3.0 (title as shown in the URL). Lawfire, Duke University. <https://sites.duke.edu/lawfire/2021/10/03/international-law-and-cyber-ops-q-a-with-mike-schmitt-about-the-status-of-tallinn-3-0/>
- [95] Cyber Law Toolkit (n.d.). Due diligence. CCDCOE. https://cyberlaw.ccdcoe.org/wiki/Due_diligence
- [96] Cooperative Cyber Defence Centre of Excellence (n.d.). Tallinn Manual page; no publication date shown. <https://ccdcoe.org/research/tallinn-manual/>
- [97] Oxford Statement on the regulation of information operations and activities (2021, June 2); Opinio Juris copy. <https://opiniojuris.org/2021/06/02/the-oxford-statement-on-international-law-protections-in-cyberspace-the-regulation-of-information-operations-and-activities/>
- [98] Global Partners Digital (n.d.). UN cyber processes: OEWG and Global Mechanism (title as shown in the URL). <https://www.gp-digital.org/un-cyber-processes-oewg-global-mechanism/>
- [99] EEAS (2026, July 21). EU statement at the first substantive session of the UN Global Mechanism on ICTs (title as shown in the URL). https://www.eeas.europa.eu/delegations/un-new-york/eu-statement-un-global-mechanism-icts-international-security-first-substantive-session-key-eu_en
- [100] Abi Chedid, D. S. (2025, October 3). Analytical Study on The Legality of Use of Sonic Boom as a 'Psywar' Tool Under International and Humanitarian Law. International Review of Law 14(2), Qatar University Press. DOI 10.29117/irl.2025.0335; read once through a summarizing tool. <https://journals.qu.edu.qa/index.php/IRL/article/view/5455>
- [101] UK Government, Intelligence Analysis (2025, March 24). Explaining Uncertainty in UK Intelligence Assessment (Professional Head of Intelligence Assessment, PHIA, Probability Yardstick); read on October 22, 2026. GOV.UK. <https://www.gov.uk/government/publications/explaining-uncertainty-in-uk-intelligence-assessment/explaining-uncertainty-in-uk-intelligence-assessment>
- [102] Cyber Law Toolkit (n.d.). Olympic Destroyer (2018). CCDCOE. [https://cyberlaw.ccdcoe.org/wiki/Olympic_Destroyer_\(2018\)](https://cyberlaw.ccdcoe.org/wiki/Olympic_Destroyer_(2018))

- [103] Council of the European Union (2026, July 13). Sanctions against cyber-attacks, policy page (review of 13 July 2026). <https://www.consilium.europa.eu/en/policies/sanctions-against-cyber-attacks/>
- [104] EEAS (2026, October 8). Russia's hybrid activities: Council prolongs restrictive measures until October 2027 (title as shown in the URL). https://www.eeas.europa.eu/delegations/ukraine/russias-hybrid-activities-council-prolongs-restrictive-measures-until-october-2027_en
- [105] Council of the European Union (2026, September 24). Decision (CFSP) 2026/2164, hybrid-threats regime listing. EUR-Lex. <https://eur-lex.europa.eu/eli/dec/2026/2164/oj/eng/pdf>
- [106] European Union (2022). Directive (EU) 2022/2555 (NIS2), recital 102. EUR-Lex. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32022L2555>
- [107] NATO (n.d.). Page on the invocation of Article 5 after 11 September 2001 (title not recorded). https://www.nato.int/cps/en/natohq/topics_110496.htm
- [108] PBS NewsHour (2025). NATO allies held Article 4 consultations after Russian drone incursion: here's what that means (title as shown in the URL); press report. <https://www.pbs.org/newshour/amp/world/nato-allies-held-article-4-consultations-after-russian-drone-incursion-heres-what-that-means>
- [109] NATO (2026, July 8). Ankara Summit Declaration; copy of the official text hosted by Affari Italiani. https://www.affaritaliani.it/wp-content/uploads/2026/07/The-Ankara-Summit-Declaration_-_NATO-Official-text.pdf
- [110] eucrim (2025). Infringement proceedings for not transposing sanctions directive (title as shown in the URL). <https://eucrim.eu/news/infringement-proceedings-for-not-transposing-sanctions-directive/>
- [111] Global Sanctions (2026, October 6). European Commission sends reasoned opinion to Ireland for failure to notify measures transposing sanctions criminalisation directive (title as shown in the URL). <https://globalsanctions.com/2026/10/european-commission-sends-reasoned-opinion-to-ireland-for-failure-to-notify-measures-transposing-sanctions-criminalisation-directive/>
- [112] European Commission (2025, May 7). Commission calls on 19 Member States to fully transpose the NIS2 Directive (title as shown in the URL). <https://digital-strategy.ec.europa.eu/en/news/commission-calls-19-member-states-fully-transpose-nis2-directive>
- [113] European Commission (2026). NIS2 transposition page; its footer date (1 July 2025) conflicts with content dated to 27 July 2026. <https://digital-strategy.ec.europa.eu/en/policies/nis-transposition>
- [114] Giannopoulos, G., Smith, H., and Theocharidou, M. (eds) (2021). The Landscape of Hybrid Threats: A Conceptual Model (Public Version). EUR 30585 EN (JRC123305), Publications Office of the European Union; text read as far as its subsection 3.2.3. https://www.hybridcoe.fi/wp-content/uploads/2021/02/conceptual_framework-reference-version-shortened-good_cover_-_publication_office.pdf
- [115] Monaghan, S. (ed.) (2019, March). Countering Hybrid Warfare. Multinational Capability Development Campaign, phase CHW2 (June 2017 to December 2018), with NUI; text read to p. 64. https://assets.publishing.service.gov.uk/media/5c8141e2e5274a2a51ac0b34/concepts_mcdc_countering_hybrid_warfare.pdf
- [116] EEAS (2026, March). 4th EEAS Threat Report on FIMI (2025 data); text read to about p. 31. https://www.eeas.europa.eu/sites/default/files/2026/documents/EEAS%204th%20Threat%20Report_web.pdf
- [117] DISARM Foundation (n.d.). DISARM framework, frequently asked questions. <https://www.disarm.foundation/faqs>

- [118] arXiv (2025). Paper listing the DISARM Red tactic stages (arXiv 2502.20997; title not recorded). <https://arxiv.org/pdf/2502.20997>
- [119] Strom, B. E., et al. (2020, March). MITRE ATT&CK: Design and Philosophy, MP180360R1. MITRE. https://attack.mitre.org/docs/ATTACK_Design_and_Philosophy_March_2020.pdf
- [120] Caltagirone, S., Pendergast, A., and Betz, C. (2013, July 5). The Diamond Model of Intrusion Analysis. Technical report ADA586960, Defense Technical Information Center; text read to section 9. <https://apps.dtic.mil/sti/pdfs/ADA586960.pdf>
- [121] Hutchins, E. M., Cloppert, M. J., and Amin, R. M. (year and venue not shown in the PDF text). White paper on intelligence-driven network defense (intrusion kill chain; title as shown in the URL). Lockheed Martin. <https://lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf>
- [122] François, C. (2019, September 20). Actors, Behaviors, Content: A Disinformation ABC. Working paper, Transatlantic High Level Working Group on Content Moderation Online and Freedom of Expression; Graphika and Berkman Klein Center. <https://democrats-science.house.gov/download/francois-addendum&download=1>
- [123] UNHCR (n.d.). Information integrity handbook, page describing the five parts actor, behaviour, content, degree and effect. <https://www.unhcr.org/handbooks/informationintegrity/node/904>
- [124] Wardle, C., and Derakhshan, H. (2017, September 27). Information Disorder: Toward an Interdisciplinary Framework for Research and Policy Making (title as shown in the URL). Council of Europe report DGI(2017)09; text read to Part 1. <https://rm.coe.int/information-disorder-toward-an-interdisciplinary-framework-for-research/168076277c>
- [125] DefenseScoop (2022, December 28). 2022 in review: DoD and Congress take aim at improving operations in the information environment (title as shown in the URL). <https://defensescoop.com/2022/12/28/2022-in-review-dod-and-congress-take-aim-at-improving-operations-in-the-information-environment/>
- [126] NATO (2023). AJP-10.1, Allied Joint Doctrine for Information Operations (paragraph numbers as recorded by the authors; promulgation date not read). UK-hosted copy. https://assets.publishing.service.gov.uk/media/650c03bf52e73c000d9425bb/AJP_10_1_Info_Ops_UK_web.pdf
- [127] NATO (n.d.). AJP-10, Allied Joint Doctrine for Strategic Communications, Change 1, para 1.18. UK-hosted copy. https://assets.publishing.service.gov.uk/media/6525459d244f8e00138e7343/AJP_10_Strat_Comm_Change_1_web.pdf
- [128] U.S. Department of Justice (2024, September 4). Release No. 24-1095 on the disruption of a covert Russian government-sponsored foreign malign influence operation (Doppelganger; title as shown in the URL). <https://www.justice.gov/archives/opa/pr/justice-department-disrupts-covert-russian-government-sponsored-foreign-malign-influence>
- [129] United States v. Certain Domains, No. 24-mj-1395 (E.D. Pa.), Doc. 4, FBI affidavit filed 2024, September 4 (allegations; the copy read stops at para 107 with no exhibits). https://www.justice.gov/d9/2024-09/doppelganger_affidavit_9.4.24.pdf
- [130] U.S. Department of the Treasury (2024, March 20). Press release jy2195, designations of Social Design Agency, Structura, Ilya Gambashidze and Nikolai Tupikin under E.O. 14024 (title not recorded). <https://home.treasury.gov/news/press-releases/jy2195>
- [131] U.S. Department of the Treasury (2024, December 31). Press release jy2766 on the designations of the Center for Geopolitical Expertise and the IRGC-linked Cognitive Design Production Center (title not recorded). <https://home.treasury.gov/news/press-releases/jy2766>

- [132] U.S. Department of Justice (2018, February 16). Grand jury indicts thirteen Russian individuals and three Russian companies for scheme to interfere (title as shown in the URL, truncated). <https://www.justice.gov/opa/pr/grand-jury-indicts-thirteen-russian-individuals-and-three-russian-companies-scheme-interfere>
- [133] United States v. Internet Research Agency LLC et al., D.D.C., Criminal Action No. 18-cr-0032-2 (DLF), Memorandum Opinion of 13 August 2018 denying Concord Management's motion to dismiss; govinfo copy. https://www.govinfo.gov/content/pkg/USCOURTS-dcd-1_18-cr-00032/pdf/USCOURTS-dcd-1_18-cr-00032-0.pdf
- [134] U.S. Department of Justice (2025, March 5). Justice Department charges 12 Chinese contract hackers and law enforcement officers in global computer intrusion campaigns (title as shown in the URL; United States v. Wu Haibo et al., S.D.N.Y.; United States v. Yin Kecheng and United States v. Zhou Shuai et al., D.D.C.). <https://www.justice.gov/opa/pr/justice-department-charges-12-chinese-contract-hackers-and-law-enforcement-officers-global>
- [135] U.S. Department of Justice, E.D.N.Y. (2023, April 17). 34 officers of the People's Republic of China national police charged with perpetrating transnational repression (title as shown in the URL); no docket number on the release. <https://www.justice.gov/usao-edny/pr/34-officers-peoples-republic-china-national-police-charged-perpetrating-transnational>
- [136] U.S. Department of Justice (2021, November 18). Two Iranian nationals charged for cyber-enabled disinformation and threat campaign designed (title as shown in the URL, truncated); no docket number on the release. <https://justice.gov/opa/pr/two-iranian-nationals-charged-cyber-enabled-disinformation-and-threat-campaign-designed>
- [137] Special Counsel's Office, U.S. Department of Justice (2019, March). Report on the Investigation into Russian Interference in the 2016 Presidential Election, vol. I (redacted version released 18 April 2019; heavily redacted excerpt read to p. 35). <https://www.justice.gov/archives/sco/file/1373816/dl>
- [138] Al Jazeera (2021, November 19). Iran condemns US sanctions over meddling in presidential vote (title as shown in the URL). <https://www.aljazeera.com/news/2021/11/19/iran-condemns-us-sanctions-over-meddling-in-presidential-vote>
- [139] Reuters via Daily Maverick (2022, November 7). Russia's Prigozhin admits interfering in US elections (title as shown in the URL). <https://www.dailymaverick.co.za/article/2022-11-07-russias-prigozhin-admits-interfering-in-us-elections.md>
- [140] Lindsay (2015). Tipping the scales: the attribution problem and the feasibility of deterrence against cyberattack. *Journal of Cybersecurity* 1(1), 53-67. DOI 10.1093/cybsec/tyv003. <https://academic.oup.com/cybersecurity/article/1/1/53/2367566>
- [141] Humprecht, E., Esser, F., and Van Aelst, P. (2020). Resilience to Online Disinformation: A Framework for Cross-National Comparative Research. *International Journal of Press/Politics* 25(3), 493-516. DOI 10.1177/1940161219900126; abstract read. <https://api.crossref.org/works/10.1177/1940161219900126>
- [142] Levin (2016). When the great power gets a vote: the effects of great power electoral interventions on election results. *International Studies Quarterly* 60(2), 189-202. DOI 10.1093/isq/sqv016; abstract-level page read. <https://hub.hku.hk/handle/10722/260221>
- [143] Council of the European Union (2022, May 10). Russian cyber operations against Ukraine: declaration by the High Representative on behalf of the European Union (title as shown in the URL). <https://www.consilium.europa.eu/en/press/press-releases/2022/05/10/russian-cyber-operations-against-ukraine-declaration-by-the-high-representative-on-behalf-of-the-european-union/>

- [144] U.S. Department of Justice (2020, October 15 return date). Indictment of six officers of GRU Military Unit 74455, W.D. Pennsylvania; criminal number illegible in the text read. <https://www.justice.gov/opa/press-release/file/1328521/download>
- [145] Goldstein, J. A., Chao, J., Grossman, S., Stamos, A., and Tomz, M. (2024). How persuasive is AI-generated propaganda? PNAS Nexus 3(2), pgae034. DOI 10.1093/pnasnexus/pgae034; Crossref record, sample size not retrieved. <https://api.crossref.org/works/10.1093/pnasnexus/pgae034>
- [146] Paul and Matthews (2016, July 11). The Russian "Firehose of Falsehood" Propaganda Model: Why It Might Work and Options to Counter It. RAND PE-198. <https://www.rand.org/pubs/perspectives/PE198.html>
- [147] Galeotti (2017, August). Controlling Chaos: How Russia Manages Its Political War in Europe. ECFR/228, European Council on Foreign Relations. https://ecfr.eu/archive/page/-/ECFR228_-_CONTROLLING_CHAOS1.pdf
- [148] The Register (2025, March 6). FBI: China pays up to USD 75k per hacked inbox (title as shown in the URL). https://www.theregister.com/2025/03/06/fbi_china_pays_75k_per/
- [149] Associated Press via Fortune (2025, March 6). 12 Chinese hackers and officials charged by US prosecutors (title as shown in the URL). <https://www.fortune.com/2025/03/06/12-chinese-hackers-officials-charged-us-prosecutors-global-cybercrime>
- [150] The Bell (2023, July 4). Prigozhin abandons media empire following failed mutiny (title as shown in the URL). <https://en.thebell.io/prigozhin-abandons-media-empire-following-failed-mutiny/>
- [151] Nimmo (2020, September). The Breakout Scale: measuring the impact of influence operations. Brookings Institution; page read, PDF not retrieved. <https://www.brookings.edu/articles/the-breakout-scale-measuring-the-impact-of-influence-operations/>
- [152] Bail et al. (2020). Study of interaction with Russian Internet Research Agency accounts and political attitudes (title not recorded). PNAS 117(1), 243-250. DOI 10.1073/pnas.1906420116; read through the Duke Scholars record and PsyPost. <https://scholars.duke.edu/publication/1422398>
- [153] Eady et al. (2023, January). Study of exposure to Russian influence accounts on Twitter in 2016 (title not recorded). Nature Communications. DOI 10.1038/s41467-022-35576-9. <https://www.nature.com/articles/s41467-022-35576-9>
- [154] Dylan, H., and Maguire, T. J. (2022). Article on the US and UK disclosure of Russian invasion plans (accepted manuscript with a different working title; title not recorded). Survival 64(4). <https://fileserv-az.core.ac.uk/download/586991258.pdf>
- [155] Santos Okholm, Ebrahimi Fard, and ten Thij (2024). Blocking the information war? Testing the effectiveness of the EU censorship of RT and Sputnik (title as shown in the URL). Internet Policy Review 13(3). <https://policyreview.info/articles/analysis/blocking-information-war-testing-effectiveness-eu-censorship>
- [156] Machtiger, P. (2020, October 29). The latest GRU indictment: a failed exercise in deterrence (title as shown in the URL); opinion piece. Just Security. <https://www.justsecurity.org/73071/the-latest-gru-indictment-a-failed-exercise-in-deterrence/>
- [157] Christian Science Monitor (2020, March 17). Why US dropping charges against Russians named by Mueller report (title as shown in the URL). <https://www.csmonitor.com/USA/Justice/2020/0317/Why-US-dropping-charges-against-Russians-named-by-Mueller-report>
- [158] TASS (2018, February 15). Kremlin reply to the NotPetya accusations (title not recorded). <https://tass.com/politics/990154/amp>

- [159] Reuters via Malay Mail (2022, May 11). Russia downed satellite internet in Ukraine, say Western officials (title as shown in the URL). <https://www.malaymail.com/news/world/2022/05/11/russia-downed-satellite-internet-in-ukraine-say-western-officials/2058256>
- [160] CyberScoop (2022). Viasat hack: Russia, UK, EU, US and Ukraine (title as shown in the URL). <https://cyberscoop.com/viasat-hack-russia-uk-eu-us-ukraine/>
- [161] TASS (2020, October 20). Peskov's reply to the British and US accusations on the Olympic Games (title not recorded). <https://tass.com/politics/1214217/amp>
- [162] Cyber Law Toolkit (n.d.). Viasat KA-SAT attack (2022). CCDCOE. [https://cyberlaw.ccdcoe.org/wiki/Viasat_KA-SAT_attack_\(2022\)](https://cyberlaw.ccdcoe.org/wiki/Viasat_KA-SAT_attack_(2022))
- [163] International Military Tribunal for the Far East (judgment; year not shown on the page read). Judgment, Chapter V, on the Mukden incident (pp. 553 to 556). HyperWar, ibiblio. <https://www.ibiblio.org/hyperwar/PTO/IMTFE/IMTFE-5.html>
- [164] Harvard Law School Nuremberg Trials Project (n.d.). Speech to the military, 22 August 1939 (document PS-1014; project summary). <https://nuremberg.law.harvard.edu/documents/452263-speech-to-the-military>
- [165] International Military Tribunal (1945, November 30). Testimony of Lahousen, IMT transcript, vol. II (mirror copy; cited by the prosecution on 20 December 1945, vol. II, p. 450). <https://phdn.org/archives/holocaust-history.org/works/imt/imtnew/02/htm/t450.htm>
- [166] Yle (2009, December 2; updated 2012, April 21). Article on the shelling of Mainila and the start of the Winter War (title not recorded). <https://yle.fi/a/3-5964904>
- [167] Finnish Military Museum (n.d.). The start of the war (Winter War information board). <https://infotaulu.museomilitaria.fi/en/winter-war/the-start-of-the-war>
- [168] Wikipedia (n.d.). Shelling of Mainila (secondary compilation, used as a pointer to references). https://en.wikipedia.org/wiki/Shelling_of_Mainila
- [169] Ynetnews (n.d.). Article on Israel's recognition of members of the Lavon Affair network on 30 March 2005 (title not recorded). <https://www.ynetnews.com/article/3065838>
- [170] Israel Hayom (2015). Article on the declassified Defense Ministry transcripts on the Lavon Affair released on 11 May 2015 (title not recorded). <https://www.israelhayom.co.il/article/281149>
- [171] German History in Documents and Images (n.d.). Decree of the Reich President for the Protection of the People and State (Reichstag Fire Decree), February 28, 1933. GHI Washington. <https://germanhistorydocs.org/en/nazi-germany-1933-1945/decreed-of-the-reich-president-for-the-protection-of-the-people-and-state-reichstag-fire-decree-february-28-1933>
- [172] History News Network (2014, March 17). Historian says Nazis were to blame for the Reichstag fire (title truncated in the URL; contains a date error for the decree). <https://www.historynewsnetwork.org/article/historian-says-nazis-were-to-blame-for-the-reichst>
- [173] McCarthy, B., and Jacobson, L. (2022, February 8). False flags: they're real, but ... (PolitiFact headline, truncated); republished by Poynter on 9 February 2022. <https://www.poynter.org/fact-checking/2022/what-is-a-false-flag/>
- [174] Mannucci Pacini and Zucchelli (2020). Strage di Piazza Fontana: indagini e processi (title as shown in the URL; write-up of a 2019 talk, no docket numbers). Sistema Penale. https://www.sistemapenale.it/pdf_contenuti/1593334874_mannuccipacinizucchelli-2020a-strage-piazz-fontana-indagini-processi.pdf

- [175] Pellegrino, G., President of the Italian Senate Commissione stragi (2001, January 9). Appunti per una relazione conclusiva (working draft; not an adopted report). Camera dei deputati, Doc. XXIII n. 64, vol. 1, t. 1. https://legislature.camera.it/_dati/leg13/lavori/doc/xxiii/064v01t01_RS/00000001.pdf
- [176] Domani (n.d.). Peteano quando l'arma depistò una strage di carabinieri (title as shown in the URL). <https://www.editorialedomani.it/peteano-quando-larma-depisto-una-strage-di-carabinieri-pr8cvq8t>
- [177] Domani (2025, December 8). La strage di Peteano e le confessioni del soldato politico Vinciguerra (title as shown in the URL). <https://www.editorialedomani.it/la-strage-di-peteano-e-le-confessioni-del-soldato-politico-vinciguerra-utscpueb>
- [178] Italian Parliament (1991, April 17). Doc. XXIII n. 32, report on Gladio (title not recorded). <https://www.parlamento.it/service/PDF/PDFServer/BGT/907543.pdf>
- [179] Italian Parliament (1991, December 4). Doc. XXIII n. 42, report on Gladio (title not recorded). <https://www.parlamento.it/service/PDF/PDFServer/BGT/909961.pdf>
- [180] Il Post (2020, October 24). Gladio e stay behind (title as shown in the URL). <https://www.ilpost.it/2020/10/24/gladio-stay-behind/>
- [181] Il Post (2023, August 7). Le sentenze sulla strage di Bologna (title as shown in the URL). <https://www.ilpost.it/2023/08/07/sentenze-strage-bologna/>
- [182] AGI (2025, January 7). Strage Bologna, Corte d'Assise, bomba, Paolo Bellini (title as shown in the URL). <https://www.agi.it/cronaca/news/2025-01-07/strage-bologna-corte-assise-bomba-paolo-bellini-29452597/>
- [183] Korkut, T. (2009, September 8). What happened on 6-7 September (title as shown in the URL). bianet. <https://bianet.org/english/minorities/116914-what-happened-on-6-7-september>
- [184] Wikipedia (n.d.). Istanbul pogrom (secondary compilation, used as a pointer to references). https://en.wikipedia.org/wiki/Istanbul_pogrom
- [185] Wikipedia (n.d.). Piazza Fontana bombing (secondary compilation, used as a pointer to references). https://en.wikipedia.org/wiki/Piazza_Fontana_bombing
- [186] Counter Terrorism Policing (2025). Group convicted after Russian-ordered arson attack in London (title as shown in the URL). <https://www.counterterrorism.police.uk/group-convicted-after-russian-ordered-arson-attack-in-london/>
- [187] Al Jazeera (2026, January 16). Lithuania charges multinational gang of Russian-directed sabotage attacks (title as shown in the URL). <https://www.aljazeera.com/news/2026/1/16/lithuania-charges-multinational-gang-of-russian-directed-sabotage-attacks>
- [188] Polish Ministry of Justice (2025, May 11). Joint statement by the Minister of Justice and Prosecutor General and the Minister of Internal Affairs and Administration, Coordinator of Secret Services (title as shown in the URL). <https://www.gov.pl/web/justice/joint-statement-by-the-minister-of-justice---prosecutor-general-and-the-minister-of-internal-affairs-and-administration-coordinator-of-secret-services>
- [189] Polish Ministry of Foreign Affairs (2025, September 10). Statement by Deputy Prime Minister Sikorski on the violation of Polish airspace (title as shown in the URL). <https://www.gov.pl/web/denmark/statement-by-deputy-prime-minister-sikorski-on-the-violation-of-polish-airspace>
- [190] Christian Science Monitor (2025, September 10). Russia, Poland drone incursion, Ukraine (title as shown in the URL). <https://www.csmonitor.com/World/Europe/2025/0910/russia-poland-drone-incursion-ukraine>

- [191] Reuters via The Star (2026, June 25). Danish police find no proof drones caused Copenhagen airport shutdown (title as shown in the URL). <https://www.thestar.com.my/news/world/2026/06/25/danish-police-find-no-proof-drones-caused-copenhagen-airport-shutdown>
- [192] CISA, NSA, FBI and partners (2024, February 7). Joint cybersecurity advisory AA24-038A on Volt Typhoon (title as shown in the URL). <https://www.cisa.gov/news-events/cybersecurity-advisories/aa24-038a>
- [193] WION, citing CNA (2026, January 11). China uses fake news sites, bot networks for cognitive warfare, NSB says (title as shown in the URL). <https://wionews.com/world/china-uses-fake-news-sites-bot-networks-for-cognitive-warfare-nsb-1768131309713>
- [194] Euronews (2026, July 3). Two Romanian men jailed by UK judge over London stabbing of Iranian TV presenter (title as shown in the URL). <https://www.euronews.com/my-europe/2026/07/03/two-romanian-men-jailed-by-uk-judge-over-london-stabbing-of-iranian-tv-presenter>
- [195] AP via PBS NewsHour (2026, July). UK says Iran-backed group behind arson and vandalism attacks on Jewish community (title as shown in the URL). <https://pbs.org/newshour/world/uk-says-iran-backed-group-behind-arson-and-vandalism-attacks-on-jewish-community>
- [196] Baker McKenzie (2026). Multiple governments issue joint alert on North Korean IT workers (title as shown in the URL; read through a tool summary on October 22, 2026, which lists the eleven countries and no agencies; the State Department text was not fetched). <https://sanctionsnews.bakermckenzie.com/multiple-governments-issue-joint-alert-on-north-korean-it-workers/>
- [197] Spanish Ministry of the Interior (2024, July 20). Tres detenidos por delitos de daños informáticos con fines terroristas (title as shown in the URL). <https://www.interior.gob.es/opencms/en/detail-pages/article/Tres-detenidos-por-delitos-de-danos-informaticos-con-fines-terroristas/>
- [198] El Español (2026, July 11). Report on the CCN-CERT assessment of cyber threats and trends (title not recorded). https://elespanol.com/espana/20260711/espana-ucrania-israel-pais-atacado-hackers-rusos-chinos-ultimo-informe-cni/1003744316174_0.amp.html
- [199] Jones, S. G. (2025, March 18). Russia's shadow war against the West. CSIS. <https://www.csis.org/analysis/russias-shadow-war-against-west>
- [200] ICCT and GLOBSEC (2026, February 23). More of the same: Russia's crime-terror nexus and criminality as a tool of hybrid warfare, revisited (title as shown in the URL). ICCT. <https://icct.nl/publication/more-same-russias-crime-terror-nexus-criminality-tool-hybrid-warfare-revisited>
- [201] EBU Spotlight (2025, September 16). Disinformation floods social media (title as shown in the URL). <https://spotlight.ebu.ch/p/disinformation-floods-social-media>
- [202] LRT (2026, March 9). GPS jamming incidents reported by pilots continue in Lithuania, regulator says (title as shown in the URL). <https://www.lrt.lt/en/news-in-english/19/2862230/gps-jamming-incidents-reported-by-pilots-continue-in-lithuania-regulator-says>
- [203] The Record (2025). Allied spy agencies blame Chinese companies for Salt Typhoon (title as shown in the URL; advisory of 27 August 2025). <https://therecord.media/allied-spy-agencies-blame-chinese-companies-salt-typhoon>
- [204] The Hacker News (2024, October). China accuses US of fabricating Volt Typhoon (title as shown in the URL). <https://thehackernews.com/2024/10/china-accuses-us-of-fabricating-volt.html>
- [205] CSO Online (2025). China alleges US cyber espionage during the Asian Winter Games, names 3 NSA agents (title as shown in the URL). <https://www.csoonline.com/article/3962756/china-alleges-us-cyber-espionage-during-the-asian-winter-games-names-3-nsa-agents.html>

- [206] The Hacker News (2026, March). Iran-linked hackers breach FBI (title as shown in the URL). <https://thehackernews.com/2026/03/iran-linked-hackers-breach-fbi.html>
- [207] Dataminr (2026, July 2). Cyber intel brief: Handala Hack confirmed Iran MOIS attribution, unverified command claims (title as shown in the URL; vendor). <https://www.dataminr.com/resources/intel-brief/cyber-intel-brief-handala-hack-confirmed-iran-mois-attribution-unverified-command-claims/>
- [208] Irish Times (2026, April 23). Iran behind 17 low-level hybrid warfare attacks in and Europe, analysts say (title as shown in the URL). <https://www.irishtimes.com/world/europe/2026/04/23/iran-behind-17-low-level-hybrid-warfare-attacks-in-and-europe-analysts-say/>
- [209] UK Government (2026, July 17). First groups designated as threat to national security (title as shown in the URL). <https://www.gov.uk/government/news/first-groups-designated-as-threat-to-national-security>
- [210] Brussels Times (2026, September 24). Iran behind attacks of Jewish sites in Liège and Antwerp, says Belgian intelligence agency (title as shown in the URL). <https://www.brusselstimes.com/2333715/iran-behind-attacks-of-jewish-sites-in-liege-and-antwerp-says-belgian-intelligence-agency>
- [211] U.S. Department of the Treasury (2026, March 12). Press release sb0416 on designations of North Korean IT-worker networks (title not recorded). <https://home.treasury.gov/news/press-releases/sb0416>
- [212] U.S. Department of Justice (2025, June 30). Justice Department announces coordinated nationwide actions to combat North Korean remote IT worker schemes (title as shown in the URL, truncated). <https://www.justice.gov/opa/pr/justice-department-announces-coordinated-nationwide-actions-combat-north-korean-remote>
- [213] WorldECR (2025, October). North Korea stole USD 2.84bn in crypto over 21 months, allied monitoring team reports (title as shown in the URL; report of the Multilateral Sanctions Monitoring Team of 22 October 2025; that report was not read). <https://www.worldecr.com/news/north-korea-stole-2-84b-in-crypto-over-21-months-allied-monitoring-team-reports/>
- [214] Rodríguez, Á. (2026, October 9). Infiltrators on the payroll: the threat of fake North Korean IT workers and the exposure of Spanish organizations. ISDO Security and Intelligence Department. <https://isdo.ch>
- [215] The Record (2025, July 16). International police takedown of NoName057(16) (title as shown in the URL). <https://therecord.media/international-police-takedown-noname-hacker>
- [216] The Register (2025, July 16). Russian hacktivist bust (title as shown in the URL). https://www.theregister.com/2025/07/16/russian_hacktivist_bust/
- [217] Anthropic (2025, November). Report on the GTG-1002 campaign (title not recorded; vendor self-report). <https://www-cdn.anthropic.com/d7dd50dd1185f59be051b307150d877f2b82bd2c.pdf>
- [218] BleepingComputer (2025). Anthropic claims of Claude AI automated cyberattacks met with doubt (title as shown in the URL). <https://www.bleepingcomputer.com/news/security/anthropic-claims-of-claude-ai-automated-cyberattacks-met-with-doubt/>
- [219] French Ministry for Europe and Foreign Affairs (2025, April 29). Russie: attribution de cyberattaques contre la France au service de renseignement militaire russe (title as shown in the URL). <https://www.diplomatie.gouv.fr/en/presse-et-ressources/decouvrir-et-informer/actualites/russie-attribution-de-cyberattaques-contre-la-france-au-service-de-renseignement-militaire-russe>

- [220] BuzzFeed News (2015, April). Experts say Russians may have posed as ISIS to hack French TV (title as shown in the URL; cites FireEye, vendor). <https://www.buzzfeednews.com/article/sheerafrenkel/experts-say-russians-may-have-posed-as-isis-to-hack-french-t>
- [221] U.S. Department of Justice (2018, July 13). Grand jury indicts 12 Russian intelligence officers for hacking offenses related to the 2016 election (title as shown in the URL; the indictment at [justice.gov/archives/sco/file/1080281](https://www.justice.gov/archives/sco/file/1080281) carries the ECF header Case 1:18-cr-00215-ABJ, D.D.C., filed July 13, 2018). <https://www.justice.gov/archives/opa/pr/grand-jury-indicts-12-russian-intelligence-officers-hacking-offenses-related-2016-election>
- [222] Wikipedia (n.d.). Guccifer 2.0 (secondary compilation, used as a pointer to references). https://en.wikipedia.org/wiki/Guccifer_2.0
- [223] CISA (2023, December 1; updated 2024, December 18). Advisory AA23-335A on an IRGC-affiliated cyber persona (title not recorded). <https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-335a>
- [224] U.S. Department of the Treasury (2024, February 2). Press release jy2072 on sanctions against six officials of the IRGC Cyber-Electronic Command (title not recorded). <https://home.treasury.gov/news/press-releases/jy2072>
- [225] NPR (2014, March 4). Putin says those aren't Russian forces in Crimea (title as shown in the URL). <https://www.npr.org/sections/thetwo-way/2014/03/04/285653335/putin-says-those-arent-russian-forces-in-crimea>
- [226] President of Russia (2014, April 17). Transcript of remarks on Crimea (title not recorded). Kremlin. <http://en.kremlin.ru/events/president/news/20796>
- [227] Dzehtsiarou, K. (2024). Ukraine v Russia (re Crimea): the European Court of Human Rights goes all in (title as shown in the URL). EJIL:Talk. <https://www.ejiltalk.org/ukraine-v-russia-re-crimea-the-european-court-of-human-rights-goes-all-in/>
- [228] Kyiv Post (date not recorded). Report on the MH17 rulings, including the Hague District Court judgment of 17 November 2022 (title not recorded). <https://www.kyivpost.com/amp/post/56002>
- [229] Crown Prosecution Service (2018, September 5). Statement on the charging decision in the Salisbury case (Sue Hemming, Director of Legal Services); Counter Terrorism Policing copy. <https://www.counterterrorism.police.uk/wp-content/uploads/2018/09/Crown-Prosecution-Service-DDPP.pdf>
- [230] Council of the European Union (2019, January 21). Decision (CFSP) 2019/86 (title not recorded). EUR-Lex. <https://eur-lex.europa.eu/eli/dec/2019/86/oj/eng>
- [231] Doughty Street Chambers (2025). Putin authorised Skripal attack which caused Dawn Sturgess death (title as shown in the URL; summary of the Dawn Sturgess Inquiry final report of 4 December 2025, which was not read). <https://www.doughtystreet.co.uk/news/putin-authorised-skripal-attack-which-caused-dawn-sturgess-death>
- [232] Al Jazeera (2026, July 2). German prosecutors allege Ukraine ordered Nord Stream pipeline attack (title as shown in the URL). <https://www.aljazeera.com/amp/news/2026/7/2/german-prosecutors-allege-ukraine-ordered-nord-stream-pipeline-attack>
- [233] beck-aktuell (2026, September 3). Prozess gegen Nord-Stream-Anschläge, Hamburg (title as shown in the URL). <https://www.beck-aktuell.de/heute-im-recht/rechtsgeschehen/prozess-gegen-nord-stream-anschlaege-hamburg-2026-09-03>
- [234] OHCHR (2022, December). Thematic report documenting civilian killings in Ukraine, including Bucha (title not recorded). United Nations Ukraine. <https://ukraine.un.org/en/download/122952/210727>

- [235] Reuters via GMA News (2022; date not recorded). Satellite images show civilian deaths in Bucha while it was in Russian hands, Maxar (title as shown in the URL). <https://www.gmanetwork.com/news/topstories/world/827517/satellite-images-show-civilian-deaths-in-bucha-while-it-was-in-russian-hands-maxar/story/>
- [236] Centre for Information Resilience (2022, April 13). Disinformation and denial: Russia's attempts to discredit open-source evidence of Bucha (title as shown in the URL). <https://www.info-res.org/eyes-on-russia/reports/disinformation-denial-russias-attempts-to-discredit-open-source-evidence-of-bucha/>
- [237] IAEA (2022; reported 2022, November 3). Release 181/2022 on inspections at three Ukrainian sites (title not recorded). <https://www.iaea.org/node/110793>
- [238] United Nations Ukraine (date not recorded). Press release on the Olenivka strike (title as shown in the URL). <https://ukraine.un.org/en/274954-olenivka-two-years-deadly-strike-killed-dozens-ukrainian-pows-continued-call-accountability>
- [239] World Nuclear News (2024, April 19). IAEA's Grossi explains why blame is not being attributed (title as shown in the URL, which is cut off). <https://world-nuclear-news.org/Articles/IAEA-s-Grossi-explains-why-blame-is-not-being-attributed>
- [240] UN Independent International Commission of Inquiry on Ukraine (2023, October 19). Report A/78/540 (title not recorded). <https://documents.un.org/doc/undoc/gen/n23/313/10/pdf/n2331310.pdf>
- [241] CNBC via NBC Bay Area (2023; date not recorded). Analysts say Russia likely staged the Kremlin drone attack it blamed on Ukraine and the West (title as shown in the URL). <https://www.nbcbayarea.com/news/business/money-report/false-flag-analysts-says-russia-likely-staged-kremlin-drone-attack-it-blamed-on-ukraine-and-the-west/3221399/>
- [242] Fox News (2023, May 24 report relaying the New York Times). Ukraine likely behind Kremlin drone attack, US officials say (title as shown in the URL). <https://www.foxnews.com/world/ukraine-likely-behind-kremlin-drone-attack-us-officials.amp>
- [243] AP via NBC San Diego (2018, April 18). Britain spy case: watchdog rejects Russia nerve agent claim (title as shown in the URL). <https://www.nbcsandiego.com/news/national-international/britain-spy-case-watchdog-rejects-russia-nerve-agent-claim/2046673/>
- [244] Dawn (2026, April 23). Report on a statement by Pakistan's information minister on the Pahalgam attack (title not recorded). <https://www.dawn.com/news/amp/1994214>
- [245] All India Radio (2025, December 16). NIA files chargesheet in Pahalgam terror attack case (title as shown in the URL). <https://www.newsonair.gov.in/nia-files-chargesheet-in-pahalgam-terror-attack-case>
- [246] 9/11 Commission (2004, July 22). Executive summary of the final report (title as shown in the URL). https://www.9-11commission.gov/report/911Report_Exec.htm
- [247] NIST (final report of 2008, November 23; FAQ page undated). Study FAQs on the World Trade Center 7 investigation (NCSTAR 1A and 1-9). <https://www.nist.gov/world-trade-center-investigation/study-faqs/wtc-7-investigation>
- [248] Sedensky, S. J., III, State's Attorney for the Judicial District of Danbury (2013, November 25). Report on the Sandy Hook Elementary School shooting (title as shown in the URL; copy hosted by the City of Seattle). <https://www.seattle.gov/documents/departments/emergency/scientificarticles/187053363-sandy-hook-final-report.pdf>
- [249] NTSB (2025, November 18). Press release on the Francis Scott Key Bridge collapse, investigation DCA24MM031 (title not recorded). <https://www.nts.gov/news/press-releases/Pages/NR20251118.aspx>

- [250] Hawaii News Now (2024, October 2). Report on the cause and origin of the Maui wildfires (title as shown in the URL). <https://www.hawaiinewsnow.com/2024/10/02/today-report-cause-origin-maui-wildfires-be-released/>
- [251] ODNI (information through August 2021). Declassified assessment on COVID-19 origins (title as shown in the URL). <https://archive.dni.gov/files/ODNI/documents/assessments/Declassified-Assessment-on-COVID-19-Origins.pdf>
- [252] WHO (2025, June 27). WHO Scientific Advisory Group issues report on origins of COVID-19 (title as shown in the URL). <https://www.who.int/news/item/27-06-2025-who-scientific-advisory-group-issues-report-on-origins-of-covid-19>
- [253] ODNI (2017, January 6). Statement on the declassified Intelligence Community Assessment (copy mirrored by GlobalSecurity; the assessment itself was not opened). <https://www.globalsecurity.org/intell/library/news/2017/intell-170106-dni01.htm>
- [254] UK Foreign Office (2019, June 14). Gulf of Oman attacks: UK statement (title as shown in the URL). <https://www.gov.uk/government/news/gulf-of-oman-attacks-uk-statement>
- [255] Wikipedia (n.d.). June 2019 Gulf of Oman incident (secondary compilation, used as a pointer to references). https://en.wikipedia.org/wiki/June_2019_Gulf_of_Oman_incident
- [256] Polskie Radio (2026, March 5). Iran's military denies firing missile at Turkey as Middle East conflict widens (title as shown in the URL). <https://www.polskieradio.pl/395/7785/artykul/3656050,irans%20military-denies-firing-missile-at-turkey-as-middle-east-conflict-widens>
- [257] Daily Sabah (2026; date not recorded). Iran denies targeting Türkiye after three intercepted missiles (title as shown in the URL). <https://www.dailysabah.com/politics/diplomacy/iran-denies-targeting-turkiye-after-3-intercepted-missiles/amp>
- [258] Al Jazeera (2026, March 23). Iran denies targeting US-UK Diego Garcia airbase, claims false flag attack (title as shown in the URL). <https://aljazeera.com/news/2026/3/23/iran-denies-targeting-us-uk-diego-garcia-airbase-claims-false-flag-attack>
- [259] TBS News (2026, October 5). False flag: Iran denies link with Flydubai incident, says claims based on fabricated evidence (title as shown in the URL). <https://www.tbsnews.net/worldbiz/middle-east/false-flag-iran-denies-link-flydubai-incident-says-claims-based-fabricated>
- [260] Israel Hayom (2026, September 30). Iran claims Israeli false flag in Flydubai flight brawl (title as shown in the URL). <https://www.israelhayom.com/2026/09/30/iran-claims-israeli-false-flag-flydubai-flight-brawl/>
- [261] Wikipedia (n.d.). Douma chemical attack (secondary compilation, used as a pointer to references). https://en.wikipedia.org/wiki/Douma_chemical_attack
- [262] OPCW (2019, March 1). Fact-Finding Mission report S/1731/2019 on the Douma incident of 7 April 2018 (title not recorded). <https://www.opcw.org/sites/default/files/documents/2019/03/s-1731-2019%28e%29.pdf>
- [263] OPCW (2023, January). OPCW releases third report of the Investigation and Identification Team (news item; title as shown in the URL; the report is dated 27 January 2023). <https://www.opcw.org/media-centre/news/2023/01/opcw-releases-third-report-investigation-and-identification-team>
- [264] Kyiv Post (2023; date not recorded). Report on Seymour Hersh's Nord Stream claim and the US denials (title not recorded). <https://www.kyivpost.com/post/12139>
- [265] Peisakhin and Rozenas (2018). Study of the Russian analog television signal in Ukraine and the 2014 elections (title not recorded). American Journal of Political Science 62(3), 535-550. DOI 10.1111/ajps.12355; abstract only. <https://api.crossref.org/works/10.1111/ajps.12355>

- [266] Enikolopov, Petrova, and Zhuravskaya (2011). Study of independent television and the 1999 Russian parliamentary elections (title not recorded). *American Economic Review* 101(7), 3253-3285. DOI 10.1257/aer.101.7.3253; abstract only. <https://api.crossref.org/works/10.1257/aer.101.7.3253>
- [267] Schiff, Schiff, and Bueno (2025; online 2024, February 20). The liar's dividend: can politicians claim misinformation to evade accountability? (title as shown in the URL). *American Political Science Review* 119(1), 71-90. <https://www.cambridge.org/core/journals/american-political-science-review/article/liars-dividend-can-politicians-claim-misinformation-to-evade-accountability/687FEE54DBD7ED0C96D72B26606AA073>
- [268] Chesney and Citron (2019). Deep fakes: a looming challenge for privacy, democracy, and national security (title as shown in the URL). *California Law Review*; volume and pages not confirmed. <https://www.californialawreview.org/print/deep-fakes-a-looming-challenge-for-privacy-democracy-and-national-security>
- [269] Goreis and Voracek (2019, February 11). A systematic review and meta-analysis of psychological research on conspiracy beliefs. *Frontiers in Psychology* 10, 205. DOI 10.3389/fpsyg.2019.00205; read through a summarizing fetch of the article page. <https://www.frontiersin.org/articles/10.3389/fpsyg.2019.00205/full>
- [270] Imhoff and Lamberty (2018, online May 30). How paranoid are conspiracy believers? *European Journal of Social Psychology* 48(7), 909-926. DOI 10.1002/ejsp.2494; metadata, and abstract through a bibliography record. <https://api.crossref.org/works/10.1002/ejsp.2494>
- [271] Freeman and Bentall (2017, online March 29). The concomitants of conspiracy concerns. *Social Psychiatry and Psychiatric Epidemiology* 52(5), 595-604. DOI 10.1007/s00127-017-1354-4; abstract and a summary of the main text. <https://link.springer.com/article/10.1007/s00127-017-1354-4>
- [272] Griffin, D. R., and Cobb, C. W. (2023, November). NATO/US false-flag attacks in Europe. *American Journal of Economics and Sociology* 82(5), 455. DOI 10.1111/ajes.12534; record and abstract read through EBSCO, full text and references not read; disputed secondary source, cited for no finding. <https://www.ebsco.com/articles/diplomacy-and-international-relations/c5c100c5-9ef4-5fa2-b36a-aa0f0bd84763/nato-us-false-flag-attacks-in-europe>
- [273] Wikipedia (n.d.). The American Journal of Economics and Sociology (secondary compilation, used as a pointer to references). https://en.wikipedia.org/wiki/The_American_Journal_of_Economics_and_Sociology
- [274] Wikipedia (n.d.). David Ray Griffin (secondary compilation, used as a pointer to references). https://en.wikipedia.org/wiki/David_Ray_Griffin
- [275] Wikipedia (n.d.). Daniele Ganser (secondary compilation, used as a pointer to references). https://en.wikipedia.org/wiki/Daniele_Ganser
- [276] Wikipedia (n.d.). U.S. Army Field Manual 30-31B (secondary compilation, used as a pointer to references). https://en.wikipedia.org/wiki/U.S._Army_Field_Manual_30-31B
- [277] The Objective (2025, April 30). Audiencia Nacional proceedings on the blackout and a possible cyberattack (title as shown in the URL). <https://theobjective.com/espana/tribunales/2025-04-30/audiencia-nacional-ciberataque-apagon-indicios/>
- [278] La Moncloa (2025, June 17). Council of Ministers press conference presenting the government committee's report on the 28 April blackout (title as shown in the URL). <https://www.lamoncloa.gob.es/consejodeministros/resumenes/paginas/2025/170625-rueda-de-prensa-ministros.aspx>

- [279] Consejo General del Poder Judicial (2026, January 12). The Audiencia Nacional judge closes the blackout investigation for lack of a minimal indication of terrorist sabotage (title as shown in the URL). <https://www.poderjudicial.es/cgpj/es/Poder-Judicial/Noticias-Judiciales/El-juez-de-la-Audiencia-Nacional-archiva-la-investigacion-por-el-apagon-general-del-pasado-el-28-de-abril-al-no-existir--un-minimo-indicio--de-sabotaje-terrorista>
- [280] ENTSO-E (2026, March 20). Final report on the 28 April 2025 Iberian blackout (title as shown in the URL). <https://www.entsoe.eu/publications/blackout/28-april-2025-iberian-blackout/>
- [281] Office of the Director of National Intelligence (2015, January 2; technical amendment 2022, January 21). Intelligence Community Directive 203, Analytic Standards; full text read for the second edition. ODNl. <https://archive.dni.gov/files/documents/ICD/ICD-203.pdf>
- [282] Central Intelligence Agency (2009, March). A Tradecraft Primer: Structured Analytic Techniques for Improving Intelligence Analysis. <https://www.cia.gov/resources/csi/static/Tradecraft-Primer-apr09.pdf>
- [283] Artner, Girven, and Bruce (2016; released 2017, January 9). RAND Research Report RR-1408 on structured analytic techniques in the US intelligence community (title not recorded). RAND Corporation. https://www.rand.org/pubs/research_reports/RR1408.html
- [284] Dhami, Belton, and Mandel (2019). Randomized test of analysis of competing hypotheses with practicing analysts (title not recorded). Applied Cognitive Psychology; read through the Strathprints record and a repository full text. <https://strathprints.strath.ac.uk/69049/>
- [285] UC Berkeley Human Rights Center and OHCHR (2020, December 2). Berkeley Protocol on Digital Open Source Investigations; full text read for the second edition. https://www.ohchr.org/sites/default/files/documents/publications/OHCHR_BerkeleyProtocol.pdf
- [286] Meduza (2022, February 18). Metadata reveals that DNR and LNR leaders recorded evacuation announcements two days in advance (title as shown in the URL). <https://meduza.io/en/news/2022/02/18/metadata-reveals-that-dnr-and-lnr-leaders-recorded-evacuation-announcements-two-days-in-advance>
- [287] SGDSN (n.d.). Les missions de VIGINUM; page also records Decree 2021-922 of 13 July 2021 and Decree 2026-70 of 11 February 2026. <https://www.sgdsn.gouv.fr/viginum/comprendre-viginum/les-missions-de-viginum>
- [288] Davis, E. G., and Grabo, C. M. (1972, April seminar remarks; approved for release 1994). Deception. CIA Studies in Intelligence. <https://cia.gov/resources/csi/static/Deception.pdf>
- [289] Infobae (2025, June 17). Government report on the 28 April blackout stresses a multifactorial origin and rules out a cyberattack (title as shown in the URL); press account, original report not read. <https://www.infobae.com/espana/2025/06/17/el-informe-del-gobierno-sobre-el-apagon-del-28-de-abril-recalca-un-origen-multifactorial-del-cero-energetico-y-descarta-un-ciberataque/>
- [290] Council of the European Union (n.d.). Hybrid threats (policy page with timeline). <https://www.consilium.europa.eu/en/policies/hybrid-threats/>
- [291] NATO (2026, updated January 29). Countering hybrid threats. https://www.nato.int/cps/en/natohq/topics_156338.htm
- [292] Council of the European Union (2025, December 15). Russian hybrid threats: Council sanctions twelve individuals and two entities over information manipulation and cyber-attacks (title as shown in the URL). <https://www.consilium.europa.eu/en/press/press-releases/2025/12/15/russian-hybrid-threats-council-sanctions-twelve-individuals-and-two-entities-over-information-manipulation-and-cyber-attacks/>

- [293] Council of the European Union (2026, April 21). Russian hybrid threats: EU lists two entities over information manipulation activities (title as shown in the URL). <https://www.consilium.europa.eu/en/press/press-releases/2026/04/21/russian-hybrid-threats-eu-lists-two-entities-over-information-manipulation-activities/>
- [294] European Court of Auditors (2021). Special Report 09/2021, Disinformation affecting the EU: tackled but not tamed. https://www.eca.europa.eu/Lists/ECADocuments/SR21_09/SR_Disinformation_EN.pdf
- [295] UK Government (2025, December 4). UK cracks down on Russian intelligence agency authorised by Putin to target Skripals (title as shown in the URL). GOV.UK. <https://www.gov.uk/government/news/uk-cracks-down-on-russian-intelligence-agency-authorised-by-putin-to-target-skipals>
- [296] SGDSN VIGINUM (page dated 2025, June 17; data to 2025, March 5). Analysis of the Russian information operating mode Storm-1516 (title as shown in the URL). defense.gouv.fr. <https://www.defense.gouv.fr/desinformation/aller-plus-loin/analyses-sgdsn/sgdsn-viginum-analyse-du-mode-operatoire-informatique-russe-storm>
- [297] EEAS (n.d.). Tackling disinformation, foreign information manipulation and interference (title as shown in the URL). https://www.eeas.europa.eu/eeas/tackling-disinformation-foreign-information-manipulation-interference_en
- [298] Roozenbeek, van der Linden, Goldberg, Rathje, and Lewandowsky (2022). Study of five inoculation videos against manipulation techniques (title not recorded). Science Advances 8(34), eabo6254. <https://www.science.org/doi/10.1126/sciadv.abo6254>
- [299] Modirrousta-Galian and Higham (2023). Reanalysis of two gamified inoculation interventions (title not recorded). Journal of Experimental Psychology: General 152(9), 2411-2437. DOI 10.1037/xge0001395. <https://doi.org/10.1037/xge0001395>
- [300] Hoes, Aitken, Zhang, Gackowski, and Wojcieszak (2024). Three survey experiments on interventions against misinformation in the United States, Poland and Hong Kong (title not recorded). Nature Human Behaviour 8(8), 1545-1553. DOI 10.1038/s41562-024-01884-x. <https://doi.org/10.1038/s41562-024-01884-x>
- [301] Carrozzini (2024, December 13). Commentary on the annulment of Romania's 2024 presidential election (title not recorded). Verfassungsblog. <https://verfassungsblog.de/?p=88320>
- [302] Csernaton and Wanless (2026, July 29). Assessing information ecosystems: how governments can get ahead of hybrid threats (title as shown in the URL). Carnegie Europe. <https://carnegieendowment.org/research/2026/07/assessing-information-ecosystems-how-governments-can-get-ahead-of-hybrid-threats>
- [303] European Union (2022). Directive (EU) 2022/2555 (NIS2), Articles 23, 28 and 29; official text read on EUR-Lex for the second edition. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32022L2555>
- [304] European Union (2022). Directive (EU) 2022/2555 (NIS2), Article 34; official text read on EUR-Lex for the second edition. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32022L2555>
- [305] NIST (2012, August). Special Publication 800-61 Revision 2 on computer security incident handling (title not recorded); withdrawn 2025, April 3 and superseded by Revision 3. <https://nvlpubs.nist.gov/nistpubs/specialpublications/nist.sp.800-61r2.pdf>
- [306] US Department of State, US Department of the Treasury, and FBI (2022, May 16). Guidance on North Korean information technology workers (title not recorded). <https://ofac.treasury.gov/media/923126/download?inline>

- [307] Collier, J. (2025, April 1). DPRK IT workers expanding scope and scale (title as shown in the URL). Google Threat Intelligence Group. <https://cloud.google.com/blog/topics/threat-intelligence/dprk-it-workers-expanding-scope-scale>
- [308] Gobierno de España (2020, November 5). Orden PCM/1030/2020 publishing the Procedimiento de actuación contra la desinformación, approved by the National Security Council on 6 October 2020 (title not recorded). Boletín Oficial del Estado, BOE-A-2020-13663. <https://www.boe.es/buscar/doc.php?id=BOE-A-2020-13663>
- [309] Departamento de Seguridad Nacional (n.d.). Foro contra las campañas de desinformación en el ámbito de la Seguridad Nacional (page; Orden PCM/541/2022, of 10 June 2022). <https://www.dsn.gob.es/foro-desinfo>
- [310] Gobierno de España (2025, March 15). Orden PJC/248/2025 on the procedure to draft the national strategy against disinformation campaigns (title not recorded). Boletín Oficial del Estado, BOE-A-2025-5151. https://www.boe.es/diario_boe/txt.php?id=BOE-A-2025-5151
- [311] Centro Criptológico Nacional (2018, June 6). Funciones del CCN (date as recorded on the page). <https://www.ccn.cni.es/index.php/es/menu-ccn-es/funciones-del-ccn>
- [312] La Moncloa (2025, January 14). Council of Ministers press conference presenting the draft Ley de Coordinación y Gobernanza de la Ciberseguridad (title as shown in the URL). <https://www.lamoncloa.gob.es/consejodeministros/resumenes/paginas/2025/140125-rueda-de-prensa-ministros.aspx>
- [313] European Commission (2026, July 8). Commission refers Ireland, Spain, France and the Netherlands to the Court of Justice for failing to transpose rules (title as shown in the URL). <https://digital-strategy.ec.europa.eu/en/news/commission-refers-ireland-spain-france-and-netherlands-court-justice-failing-transpose-rules>
- [314] Council of the European Union (2025, May 20). Russian hybrid threats: EU lists further 21 individuals and 6 entities and introduces sectoral measures (title as shown in the URL; page not opened). <https://www.consilium.europa.eu/en/press/press-releases/2025/05/20/russian-hybrid-threats-eu-lists-further-21-individuals-and-6-entities-and-introduces-sectoral-measures-in-response-to-destabilising-activities-against-the-eu-its-member-states-and-international-partners/>
- [315] EEAS (2025, July 15). Russian hybrid threats: EU lists nine individuals and six entities responsible for destabilising actions (title as shown in the URL). https://www.eeas.europa.eu/delegations/ukraine/russian-hybrid-threats-eu-lists-nine-individuals-and-six-entities-responsible-destabilising-actions_en
- [316] Council of the European Union (2025, October 3). Russian hybrid threats: Council prolongs restrictive measures by another year (title as shown in the URL; page not opened). <https://www.consilium.europa.eu/en/press/press-releases/2025/10/03/russian-hybrid-threats-council-prolongs-restrictive-measures-by-another-year/>
- [317] Council of the European Union (2026, January 29). Russian hybrid threats: Council sanctions six individuals over information manipulation activities (title as shown in the URL); copy hosted by the Cyprus Government. <https://www.cgc.org.cy/en/russian-hybrid-threats-council-sanctions-six-individuals-over-information-manipulation-activities/>
- [318] EEAS (2026, March 16). Russian hybrid threats: four individuals added to EU sanctions list for information manipulation activities (title as shown in the URL); Decision (CFSP) 2026/646. https://www.eeas.europa.eu/delegations/ukraine/russian-hybrid-threats-four-individuals-added-eu-sanctions-list-information-manipulation-activities_en

- [319] Eunews (2026, October 8). EU extends sanctions against Russia over hybrid activities by one year (title as shown in the URL); press report. <https://www.eunews.it/en/2026/10/08/eu-extends-sanctions-against-russia-over-hybrid-activities-by-one-year/>
- [320] Lovguiden (2026, September 24). EU sanctions former RT France head Xenia Fedorova (title as shown in the URL); Danish legal-information summary of the Council act, so secondary. <https://lovguiden.dk/en/det-offentlige/eur-lex-lovgivning/2026-09-24-eu-sanktioner-tidligere-rt-france-chef-xenia-fedorova>
- [321] Bajarunas, E. (2025, December 2). The Hybrid Threat Imperative: deterring Russia before it is too late (title as shown in the URL). CEPA; the read stopped at about 100,000 of 107,800 characters; the URL is percent-encoded at one letter to avoid a house-style vetoed word in its path. <https://cepa.org/compre%68ensive-reports/the-hybrid-threat-imperative-deterring-russia-before-it-is-too-late/>
- [322] Nelson and Sanchez (2025, May 14). Alone we stand: how Europe can counter hybrid threats in a post-transatlantic era (title as shown in the URL). ECFR. <https://ecfr.eu/article/alone-we-stand-how-europe-can-counter-hybrid-threats-in-a-post-transatlantic-era/>
- [323] European Parliamentary Research Service (2026, January 15). The European Democracy Shield: an overview. <https://epthinktank.eu/2026/01/15/the-european-democracy-shield-an-overview/>
- [324] NATO (2025, June 25). Hague Summit Declaration. https://www.nato.int/cps/en/natohq/official_texts_236705.htm
- [325] Wikipedia (n.d.). Catalangate (tertiary source relaying the Citizen Lab report of 18 April 2022). <https://en.wikipedia.org/wiki/Catalangate>
- [326] OpenAI (2024, May 1). Disrupting malicious uses of AI: Spamouflage (title as shown in the URL; vendor report). <https://openai.com/fr-FR/index/disrupting-malicious-uses-of-ai-spamouflage/>
- [327] LRT (2026, March 6). Five men charged in Lithuania over parcel arson plot linked to Russian intelligence (title as shown in the URL). <https://www.lrt.lt/en/news-in-english/19/2859285/five-men-charged-in-lithuania-over-parcel-arson-plot-linked-to-russian-intelligence>
- [328] U.S. Department of Justice (2024, September 27). Three IRGC cyber actors indicted for hack-and-leak operation designed to influence the 2024 US election (title as shown in the URL). <https://www.justice.gov/archives/opa/pr/three-irgc-cyber-actors-indicted-hack-and-leak-operation-designed-influence-2024-us>
- [329] The Local (2025, October 30). German-Russian jailed for six years in Munich sabotage trial (title as shown in the URL). <https://www.thelocal.de/20251030/german-russian-jailed-for-six-years-in-munich-sabotage-trial>
- [330] AP via News4Jax (2025, December 12). Germany summons Russian ambassador over alleged sabotage, cyberattacks and election interference (title as shown in the URL). <https://www.news4jax.com/business/2025/12/12/germany-summons-russian-ambassador-over-alleged-sabotage-cyberattacks-and-election-interference/>
- [331] Maritime Executive (2025). Taiwan jails Chinese master for causing damage to subsea cable (title as shown in the URL). <https://maritime-executive.com/article/taiwan-jails-chinese-master-for-causing-damage-to-subsea-cable>
- [332] UK National Cyber Security Centre (2025, July 18). UK call-out of Russian military intelligence use of espionage tool (title as shown in the URL). <https://www.ncsc.gov.uk/news/uk-call-out-russian-military-intelligence-use-espionage-tool>
- [333] Anadolu Agency (2025, November). Belgian authorities suspect Russian involvement in drone incidents (title as shown in the URL). <https://aa.com.tr/en/europe/belgian-authorities-suspect-russian-involvement-in-drone-incidents/3735639>

- [334] Notes from Poland (2025, November 24). Poland charges Ukrainian with assisting in Russian rail sabotage (title as shown in the URL). <https://notesfrompoland.com/2025/11/24/poland-charges-ukrainian-with-assisting-in-russian-rail-sabotage/>
- [335] Al Jazeera (2025, December 31). Police investigating damage to undersea telecom cable in Gulf of Finland (title as shown in the URL). <https://www.aljazeera.com/news/2025/12/31/police-investigating-damage-to-undersea-telecom-cable-in-gulf-of-finland>
- [336] Wikipedia (n.d.). Fitburg (tertiary source). <https://en.wikipedia.org/wiki/Fitburg>
- [337] Banque des Territoires (date not recorded). VIGINUM confirms two foreign digital interferences ahead of the municipal elections (title as shown in the URL). <https://www.banquedesterritoires.fr/viginum-confirme-deux-ingerences-numeriques-etrangeres-laproche-des-municipales>
- [338] International Crisis Group (2026, September 1). Strait of Hormuz, Iran-US/Israel trigger list, flashpoint entry. <https://www.crisisgroup.org/trigger-list/iran-usisrael-trigger-list/flashpoints/strait-hormuz>
- [339] Wikipedia (n.d.). Reichstag fire (secondary compilation, used as a pointer to references). https://en.wikipedia.org/wiki/Reichstag_fire
- [340] Gasiorowski, M. J. (2000, September). The CIA looks back at the 1953 coup in Iran (title as shown in the URL; summary of the CIA's internal history). MERIP. <https://merip.org/2000/09/the-cia-looks-back-at-the-1953-coup-in-iran/>
- [341] National Security Archive (2013). Posting of 21 CIA and 14 British documents on the 1953 coup in Iran, NSA Archive Electronic Briefing Book 435. George Washington University. <https://nsarchive2.gwu.edu/NSAEBB/NSAEBB435/>
- [342] Wikipedia (n.d.). Lavon Affair (secondary compilation, used as a pointer to references). https://en.wikipedia.org/wiki/Lavon_Affair
- [343] Oren (n.d.). A Truly Unfortunate Business (account of the Lavon Affair; first name and date not recorded; hosted at michaeloren.org). <https://michaeloren.org/clarity/a-truly-unfortunate-business/>
- [344] Shlaim, A. (1997). The Protocol of Sèvres, 1956: Anatomy of a War Plot. International Affairs 73, 509-530 (pages from a library catalogue record; author's online copy read). <https://users.ox.ac.uk/%7Essfc0005/The%20Protocol%20of%20Sevres%201956%20Anatomy%20of%20a%20War%20Plot.html>
- [345] European Parliament (1990, November 22). Resolution on Gladio, reproduced by Statewatch (Official Journal reference not retrieved). <https://www.statewatch.org/statewatch-database/gladio-resolution-of-the-european-parliament>
- [346] AGI (2025, July 1). Strage Bologna, Cassazione conferma ergastolo per Bellini (title as shown in the URL). <https://www.agi.it/cronaca/news/2025-07-01/strage-bologna-cassazione-conferma-ergastolo-per-bellini-32137995/>
- [347] Radio Free Europe/Radio Liberty (2004, January 12). Report on the closed trial for the Moscow and Volgodonsk bombings (title not recorded). <https://www.rferl.org/a/1143075.html>
- [348] Radio Free Europe/Radio Liberty (n.d.). Who blew up Moscow apartment buildings, 1999 (title as shown in the URL). <https://www.rferl.org/amp/who-blew-up-moscow-apartment-buildings-1999/31033208.html>
- [349] U.S. Senate Committee on Foreign Relations, minority staff (2018, January). Report on Russian influence, Senate Print 115-21 (title not recorded). GovInfo. <https://www.govinfo.gov/content/pkg/CPRT-115SPRT28110/html/CPRT-115SPRT28110.htm>

- [350] Independent International Fact-Finding Mission on the Conflict in Georgia (2009, September). Report, Volume I (Tagliavini Report); page numbers inferred from extracted text. Max Planck Institute for Comparative Public Law and International Law copy. https://www.mpil.de/files/pdf4/IIFFMCG_Volume_I2.pdf
- [351] European Court of Human Rights, Grand Chamber (2021, January 21). Georgia v. Russia (II), application no. 38263/08, judgment; HUDOC text as hosted by lagen.nu. <https://lagen.nu/dom/echr/001-207757>
- [352] International Criminal Court, Office of the Prosecutor (2022, December 16). Prosecutor Karim AA Khan KC announces conclusion of the investigation phase in Georgia (title as shown in the URL; read through a tool summary). <https://www.icc-cpi.int/news/prosecutor-international-criminal-court-karim-aa-khan-kc-announces-conclusion-investigation>
- [353] Council of Europe Committee of Ministers (2021, May 11). Decision on Ukraine v. Russia (re Crimea), applications nos. 20958/14 and 38334/18, recording the Grand Chamber admissibility decision of 14 January 2021. <https://rm.coe.int/0900001680a2471d>
- [354] Bellingcat (2022, February 23). Documenting and debunking dubious footage from Ukraine's frontlines (title as shown in the URL). <https://bellingcat.com/news/2022/02/23/documenting-and-debunking-dubious-footage-from-ukraines-frontlines>
- [355] Euromaidan Press (2022, February 3). Russia planned graphic video with corpses as pretext for Ukraine invasion, US intel says (title as shown in the URL). <https://euromaidanpress.com/2022/02/03/russia-planned-graphic-video-with-corporse-as-pretext-for-ukraine-invasion-us-intel-says/>
- [356] CBC (2022, February 17; report on Secretary of State Blinken's address to the UN Security Council; title not recorded). <https://www.cbc.ca/lite/story/1.6354905>
- [357] ADN/AP (2026, March 12). Russian court convicts 19 people over deadly 2024 Moscow concert hall attack (title as shown in the URL). <https://www.adn.com/nation-world/2026/03/12/russian-court-convicts-19-people-over-deadly-2024-moscow-concert-hall-attack/>
- [358] UN News (2013, September 16). Report of the UN Mission to Investigate Alleged Uses of Chemical Weapons on Ghouta (title not recorded). <https://news.un.org/en/story/2013/09/449052>
- [359] Arms Control Association (2017, November). Panel cites Syrian regime in sarin attack (title as shown in the URL; report of the OPCW-UN Joint Investigative Mechanism of 26 October 2017). <https://www.armscontrol.org/act/2017-11/news-briefs/panel-cites-syrian-regime-sarin-attack>
- [360] Federal Court of Justice (Germany), 3rd Criminal Senate (2025, December 10). Order StB 60/25 (Nord Stream), headnotes only; HRRS 2/2026. <https://www.hrr-strafrecht.de/hrr/archiv/26-02/hrrs-2-26.pdf>
- [361] NBC News (2019, June 6). Oil tanker sabotage likely work of state actor, investigation concludes (title as shown in the URL). <https://www.nbcnews.com/news/world/oil-tanker-sabotage-likely-work-state-actor-investigation-concludes-n1014951>
- [362] Euronews (2026, October 9). Iran's IRGC strikes tanker and threatens to target vessels beyond Hormuz (title as shown in the URL). <https://euronews.com/2026/10/09/irans-irgc-strikes-tanker-and-threatens-to-target-vessels-beyond-hormuz>
- [363] Al Jazeera (2026, March 31). Drone attack sparks fire on Kuwaiti tanker in UAE amid Iran's Gulf attacks (title as shown in the URL). <https://www.aljazeera.com/news/2026/3/31/drone-attack-sparks-fire-on-kuwaiti-tanker-in-uae-amid-irans-gulf-attacks>
- [364] Ship & Bunker (2026). Palau rejects reports linking its flag to tanker attacked off Oman (title as shown in the URL). <https://shipandbunker.com/news/world/515565-palau-rejects-reports-linking-its-flag-to-tanker-attacked-off-oman>

- [365] U.S. Central Command (2026, June 8). US forces disable non-compliant oil tanker in Gulf of Oman (title as shown in the URL; public release). <https://www.centcom.mil/MEDIA/PUBLIC-RELEASES/Article/4511257/us-forces-disable-non-compliant-oil-tanker-in-gulf-of-oman/>
- [366] Shafaq News (2026, July 30). Houthis deny involvement in attack on vessels in Egypt (title as shown in the URL). <https://shafaq.com/en/Middle-East/Houthis-deny-involvement-in-attack-on-vessels-in-Egypt>
- [367] Maritime Executive (2026, August 13). Egypt disavows association with cargo ship attacked by Houthis (title as shown in the URL). <https://maritime-executive.com/article/egypt-disavows-association-with-cargo-ship-attacked-by-houthis>
- [368] Maritime Executive (2026). CENTCOM says Iran attempts to intimidate shipping with false reports (title as shown in the URL; relays an interested source). <https://maritime-executive.com/article/centcom-says-iran-attempts-to-intimidate-shipping-with-false-reports>
- [369] Parliament of Australia, Joint Standing Committee on Foreign Affairs, Defence and Trade (n.d.). Report on the loss of HMAS Sydney, chapter 4 (report title and date not shown in the excerpt read; committee identified from the URL path). https://aph.gov.au/~media/wopapub/house/committee/jfadt/Sydney/sydch4_pdf.ashx
- [370] Wikipedia (n.d.). Q-ship. <https://en.wikipedia.org/wiki/Q-ship>
- [371] Wikipedia (n.d.). Perfidy. <https://en.wikipedia.org/wiki/Perfidy>
- [372] NPR (2025, October 15). Supreme Court, Alex Jones defamation judgment (title as shown in the URL). <https://www.npr.org/2025/10/15/nx-s1-5574929/supreme-court-alex-jones-defamation-judgment>
- [373] International Court of Justice (n.d.). Oil Platforms (Iran v. United States), case page, ICJ no. 90; judgment of 6 November 2003. <https://www.icj-cij.org/case/90>
- [374] The EastAfrican (2022, September 12). Uganda pays first USD 65m to DRC as war reparations (title as shown in the URL). <https://www.theeastafrican.co.ke/tea/news/east-africa/uganda-pays-first-usd65m-to-drc-as-war-reparations-3945720>
- [375] American Society of International Law (2022). ICJ delivers order in Ukraine v. Russia case (title as shown in the URL; order of 16 March 2022). International Law in Brief. <https://www.asil.org/ILIB/icj-delivers-order-ukraine-v-russia-case>
- [376] Novaya Gazeta Europe (2026, March 12). Moscow court hands perpetrators of Crocus City Hall terror attack life sentences (title as shown in the URL). <https://new.novayagazeta.eu/articles/2026/03/12/moscow-court-hands-perpetrators-of-crocus-city-hall-terror-attack-life-sentences-en-news>
- [377] International Commission of Jurists (2025, July 29). Press statement on the French al-Assad case (title as shown in the URL). <https://www.icj.org/wp-content/uploads/2025/07/20250729-FINAL-Press-statement-Syria-French-al-Assad-case.pdf>
- [378] GlobalSpec (n.d.). Catalogue record of NATO AJP-3.10, Allied Joint Doctrine for Information Operations: publication shown as inactive, 82 pages, document history dated 23 November 2009 and 7 December 2015. Standards catalogue. <https://standards.globalspec.com/std/1247508/ajp-3-10>
- [379] UK Ministry of Defence (2024, October 11). Allied Joint Doctrine for Military Public Affairs (AJP-10.3), Edition A Version 1; publication page, body of the text not read. GOV.UK. <https://www.gov.uk/government/publications/allied-joint-doctrine-for-military-public-affairs-ajp-103>

- [380] NATO (2024, October 18). NATO's approach to countering information threats. Official text. https://www.nato.int/cps/en/natohq/official_texts_231905.htm
- [381] Blatny, J. M., and Søndergaard, S. (2025). Cognitive Warfare. NATO Science and Technology Organization, Chief Scientist's report (STO-OCS-001); the PDF disclaims representing NATO's official position. <https://www.sto.nato.int/document/cognitive-warfare/>
- [382] Reding, A., Weed, K., and Ghez, J. (2010, November 12). NATO's Strategic Communications concept and its relevance for France. RAND Europe technical report TR-855/2-MOD/FR. https://www.rand.org/pubs/technical_reports/TR855z2.html
- [383] Paul, C., and others (2018, June 7). Lessons from Others for Future U.S. Army Operations in and Through the Information Environment. RAND Research Report RR-1925/1-A. https://www.rand.org/pubs/research_reports/RR1925z1.html
- [384] Aziz, M. H. R. bin, and Zulkifli, N. (2025, June). Hybrid Warfare Case Study: Syria Civil War, 2011 and Annexation of Crimea by Russia, 2014. International Journal of Social Science and Human Research 8(6), 4764-4773. DOI 10.47191/ijsshr/v8-i6-81; abstract read, PDF through a summarizing tool. <https://ijsshr.in/v8i6/Doc/81.pdf>
- [385] Jamil, H., Fatima, N., and Shahbaz, K. (2025, January-March). The Rise of Asymmetrical Warfare in Great Power Competition. Pakistan Social Sciences Review 9(1), 961-972. DOI 10.35484/pssr.2025(9-1)73; abstract read. <https://ojs.pssr.org.pk/journal/article/view/1206>
- [386] Qureshi, H. I. A., and Khalid, I. (2024, October-December). Hybrid Warfare in the 21st Century: Implications for Pakistan. Journal of Development and Social Sciences 5(4), 459-468. DOI 10.47205/jdss.2024(5-IV)39; abstract read. <https://ojs.jdss.org.pk/journal/article/view/1266>
- [387] Prasetyo, H., Suharto, P., Santoso, T. I., and Al-Mujaddid, S. (2025, August 28). Hybrid warfare and total defense strategy: Indonesia's security framework in the Indo-Pacific era (title abridged here by one adjective). International Journal of Humanities Education and Social Sciences 5(1), 311-321. DOI 10.55227/ijhess.v5i1.1786; abstract read, PDF through a summarizing tool. <https://ijhess.com/index.php/ijhess/article/view/1786>
- [388] Panngam, P. (2024, January-April). Superpowers and Hybrid Warfare. Nimit Mai Review (Thailand) 7(1), 57-65; the likely match for a list entry that gave no title. <https://so04.tci-thaijo.org/index.php/nmrj/article/download/255245/183351/1099761>
- [389] Turns, D. (2023, March 31). Ukraine war: 'false flag' operations. Cranfield University press release PR-CDS-23-37, republished from The Conversation; not peer reviewed. <https://www.cranfield.ac.uk/press/news-2023/ukraine-war-false-flag-operations>
- [390] Gambetta, D. (2024). The Red Brigade Signature: Mimic-proofing Claims of Political Violent Actions, Italy, 1969-1980. In Gambetta, D., and Hegghammer, T. (Eds.), Fight, Flight, Mimic: Identity Mimicry in Conflict. Oxford University Press. DOI 10.1093/9780191802454.003.0005; chapter abstract only partly visible. <https://doi.org/10.1093/9780191802454.003.0005>
- [391] McCarthy, B., and Jacobson, L. (2022, February 8). False flags: they're real, but ... (PolitiFact headline, truncated). PolitiFact; republished by Poynter on 9 February 2022. <https://politifact.com/article/2022/feb/08/false-flags-theyre-real-far-less-widespread-social/>
- [392] Contractor study for the Office of Net Assessment, Office of the Secretary of Defense (2013, May). China: The Three Warfares. Released in litigation with authors redacted (the Washington Free Beacon, 2014, names Stefan Halper as director); about 100,000 of 126,000 characters read. https://www.esd.whs.mil/Portals/54/Documents/FOID/Reading%20Room/Litigation_Release/Litigation%20Release%20-%20China-%20The%20Three%20Warfares%20%20201305.pdf

- [393] Gershaneck, K. K. (2024). China's "Second Battlefield": Political Warfare in Combat Operations. *Journal of Advanced Military Studies* 15(2). DOI 10.21140/mcu.j.20241502009. <https://www.usmcu.edu/Outreach/Marine-Corps-University-Press/MCU-Journal/JAMS-vol-15-no-2/Chinas-Second-Battlefield/>
- [394] Gershaneck, K. K., and Chan, E. (2025, April). Political Warfare against Intervention Forces. *Journal of Indo-Pacific Affairs*, Air University Press (volume and issue not shown; the PDF path carries 28 April 2025). <https://www.airuniversity.af.edu/JIPA/Display/Article/4167178/political-warfare-against-intervention-forces/>
- [395] Sinaga, T. M. (2025, November issue; online 2026, February 7). Grey Zone Escalation: How China's Hybrid Warfare Is Reshaping Taiwan's Security. *Insignia: Journal of International Relations* 12(2), 183-200. DOI 10.20884/1.ins.2025.12.2.19396; abstract read. <https://jos.unsoed.ac.id/index.php/insignia/article/view/19396>
- [396] Eggen, K.-A. (2025, September 19 online). A Strategy for the Weak: The Role of Information Confrontation in Russia's Grand Strategy. *Defence Studies* 26(2), 211-235. DOI 10.1080/14702436.2025.2561639; abstract read through OpenAlex, full text not read. <https://api.crossref.org/works/10.1080/14702436.2025.2561639>
- [397] Vykhodets, R. S. (2025). The Policy of Ensuring the Information and Psychological Security of the CSTO (title abridged here; the published title ends "in the Context of Global Strategic Competition"). *Vestnik RUDN. International Relations* 25(1), 78-86. DOI 10.22363/2313-0660-2025-25-1-78-86; abstract read. https://journals.rudn.ru/international-relations/article/view/43460/en_US
- [398] Stancu, M. N., and Căzănaru, G. (2024, September). Psychological Operations During the Russian War of Aggression in Ukraine. *Land Forces Academy Review* 29(3), 351-357. DOI 10.2478/raft-2024-0037; abstract read. <https://reference-global.com/article/10.2478/raft-2024-0037>
- [399] Hurkovskyi, V., Ilnytskyi, S., and Marutian, R. (2026, June 30). Transformation of Russia's Strategic Information Operations Against Ukraine (2024-2025): The Technological Dimension and Proactive Response Mechanisms. *Visnyk of Taras Shevchenko National University of Kyiv: Military-Special Sciences* 2(66), 97-103. DOI 10.17721/1728-2217.2026.66.97-103; abstract read through Crossref. <https://api.crossref.org/works/10.17721/1728-2217.2026.66.97-103>
- [400] Sanaei (2019, November). Time is of the essence: The causal effect of duration on support for war. *Journal of Peace Research* 56(6), 783-796. DOI 10.1177/0022343319834984; record page read, sample not shown. <https://ideas.repec.org/a/sae/joupea/v56y2019i6p783-796.html>
- [401] Felter, J. H., and Shapiro, J. N. (2017, Winter). Limiting Civilian Casualties as Part of a Winning Strategy: The Case for Courageous Restraint. *Daedalus* (online 9 January 2017); analytical essay, volume, pages and DOI not read. <https://www.amacad.org/publication/daedalus/limiting-civilian-casualties-part-winning-strategy-case-courageous-restraint>
- [402] Feaver, Gelpi, and Reifler (2005/06, Winter). Success Matters: Casualty Sensitivity and the War in Iraq. *International Security* (volume and pages not confirmed); abstract-level page read. <https://www.belfercenter.org/publication/success-matters-casualty-sensitivity-and-war-iraq>
- [403] Nelson, A., Rekhi, S., Souppaya, M., and Scarfone, K. (2025, April). Incident Response Recommendations and Considerations for Cybersecurity Risk Management: A CSF 2.0 Community Profile (NIST Special Publication 800-61 Revision 3, which supersedes Revision 2). DOI 10.6028/NIST.SP.800-61r3. NIST. <https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r3.pdf>
- [404] Brezinski, D., and Killalea, T. (2002, February). Guidelines for Evidence Collection and Archiving (RFC 3227, BCP 55). Internet Engineering Task Force. <https://www.rfc-editor.org/rfc/rfc3227.txt>

- [405] Kent, K., Chevalier, S., Grance, T., and Dang, H. (2006, August). Guide to Integrating Forensic Techniques into Incident Response (NIST Special Publication 800-86). NIST. <https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-86.pdf>
- [406] International Organization for Standardization and International Electrotechnical Commission (2012). ISO/IEC 27037:2012, Information technology, Security techniques, Guidelines for identification, collection, acquisition and preservation of digital evidence; catalogue page read, and only the informative scope and definitions are public. ISO. <https://www.iso.org/standard/44381.html>
- [407] European Union (2016). Regulation (EU) 2016/679 (General Data Protection Regulation), Articles 5, 33 and 34; official text, Official Journal L 119, 4 May 2016. EUR-Lex. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679>
- [408] National Institute of Standards and Technology (2022, December 15). NIST Retires SHA-1 Cryptographic Algorithm (news release). NIST. <https://www.nist.gov/news-events/news/2022/12/nist-retires-sha-1-cryptographic-algorithm>
- [409] National Security Agency, Australian Signals Directorate's Australian Cyber Security Centre, Canadian Centre for Cyber Security, and United Kingdom National Cyber Security Centre (2025, January). Content Credentials: Strengthening Multimedia Integrity in the Generative AI Era (Cybersecurity Information Sheet U/OO/109191-25, PP-25-0336, version 1.0). NSA. <https://media.defense.gov/2025/Jan/29/2003634788/-1/-1/0/CSI-CONTENT-CREDENTIALS.PDF>
- [410] Hollenbeck, S., and Newton, A. (2021, June). JSON Responses for the Registration Data Access Protocol (RDAP) (RFC 9083, STD 95). Internet Engineering Task Force. <https://www.rfc-editor.org/rfc/rfc9083.txt>
- [411] Laurie, B., Messeri, E., and Stradling, R. (2021, December). Certificate Transparency Version 2.0 (RFC 9162, Experimental). Internet Engineering Task Force. <https://www.rfc-editor.org/rfc/rfc9162.txt>
- [412] Marcoci, A., Burgman, M., Kruger, A., Silver, E., McBride, M., Singleton Thorn, F., Fraser, H., Wintle, B. C., Fidler, F., and Vercammen, A. (2019, January 7). Better Together: Reliable Application of the Post-9/11 and Post-Iraq US Intelligence Tradecraft Standards Requires Collective Analysis. *Frontiers in Psychology* 9, article 2634. DOI 10.3389/fpsyg.2018.02634. <https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2018.02634/full>
- [413] European Union (2023, July 12). Regulation (EU) 2023/1543 on European Production Orders and European Preservation Orders for electronic evidence in criminal proceedings, Articles 4, 6, 11 and 34; official text, Official Journal L 191, 28 July 2023, p. 118, applicable from 18 August 2026. EUR-Lex. <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32023R1543>
- [414] Ministerio de Gracia y Justicia (1882, September 14). Real Decreto de 14 de septiembre de 1882 por el que se aprueba la Ley de Enjuiciamiento Criminal, article 588 octies (order to preserve data); consolidated text, which carries no official value. Boletín Oficial del Estado, BOE-A-1882-6036. <https://www.boe.es/eli/es/rd/1882/09/14/1/con>
- [415] Council of Europe (2001, November 23). Convention on Cybercrime (ETS No. 185), Articles 16, 29 and 35; official text read, and the status of each State as a party not checked. Council of Europe. <https://rm.coe.int/1680081561>
- [416] Consejo Nacional de Ciberseguridad (2020, February 21). Guía nacional de notificación y gestión de ciberincidentes (national guide on cyber-incident notification and management, in Spanish), published by INCIBE; it predates NIS2 and rests on Royal Decree-Law 12/2018. INCIBE. https://www.incibe.es/sites/default/files/contenidos/guias/doc/guia_nacional_notificacion_gestion_ciberincidentes.pdf

- [417] Council of Europe, Department for the Execution of Judgments (consulted 2026, October 22). Register of just satisfaction concerning the Russian Federation (title as shown on the page; Georgia v. Russia (II), no. 38263/08, just satisfaction judgment of 28 April 2023; read through a tool summary). Council of Europe. <https://coe.int/en/web/execution/register>
- [418] European Court of Human Rights (2025, March 6). Inadmissibility decision concerning Romania, Călin Georgescu v. Romania (page title as shown; the page gives no application number; read through a tool summary). ECtHR. <https://www.echr.coe.int/w/inadmissibility-decision-concerning-romania-1>
- [419] European Court of Human Rights (2025, March 6). Press release on Călin Georgescu v. Romania, application no. 37327/24, committee decision; Romanian-language copy hosted by Antena 3 (read through a tool summary). <https://antena3.ro/pictures/documents/2025/03/06/1570-Calin-Georgescu-v.-Romania-Curtea-Europeana-a-Drepturilor-Omului-respinge-cererea-lui-Calin-Georgescu-privind-anularea-alegerilor-prezidentiale-din-Rom%C3%A2nia-1.pdf>
- [420] IRIS Merlin, European Audiovisual Observatory (n.d.). [RO] Annulment of the electoral process for the election of the President of Romania in 2024 (note on Decision no. 32 of 6 December 2024 of the Constitutional Court, file no. 3771F/2024; read through a tool summary). <https://merlin.obs.coe.int/article/10222>
- [421] Crown Prosecution Service (2025, October 24). Page on the Leyton warehouse arson prosecution (headline: "UPDATED WITH SENTENCE: How the CPS used new National Security Act legislation to prosecute the plot to sabotage Ukrainian aid warehouses on UK soil"; read through a tool summary). CPS. <https://www.cps.gov.uk/node/27578>
- [422] Appellate Court of Connecticut (2024, December 10). Lafferty v. Jones, 229 Conn. App. 487, 327 A.3d 941 (AC 46131, AC 46132, AC 46133); excerpt only, the holding was not read. Leagle. <https://leagle.com/decision/inctco20241210070>
- [423] Columbia University, Global Freedom of Expression (n.d.). Lafferty v. Jones, case summary (prints the trial docket as X06-UWY-CV18-604637-S, unconfirmed; reports the Appellate Court result of 10 December 2024; read through a tool summary). <https://globalfreedomofexpression.columbia.edu/cases/lafferty-v-jones/>
- [424] International Criminal Court, Pre-Trial Chamber I (2022, June 30). Public redacted version of the arrest warrant for Gamlet Guchmazov, ICC-01/15-41-Red, Situation in Georgia (read through a tool summary). https://icc-cpi.int/sites/default/files/CourtRecords/CR2022_05215.PDF
- [425] International Criminal Tribunal for the former Yugoslavia, Appeals Chamber (1999, July 15). Prosecutor v. Tadić, case no. IT-94-1-A, transcript of the judgment hearing (read through a tool summary). <https://icty.org/x/cases/tadic/trans/en/990715it.htm>
- [426] *bijzonderstrafrecht.nl* (n.d.). De vonnissen in de strafzaak MH17 in internationaal strafrechtelijk perspectief (title as shown; it records four judgments of 17 November 2022 and prints no ECLI; read through a tool summary). <https://www.bijzonderstrafrecht.nl/home/de-vonnissen-in-de-strafzaak-mh17-in-internationaal-strafrechtelijk-perspectief>
- [427] *ms-aktuell* (date not recorded). Nord Stream: Neue Details zur Anklage vor dem Prozess in Hamburg (gives the file number 3 St 2/26 and 16 hearing dates to 17 December 2026; read through a tool summary). <https://ms-aktuell.de/?p=166427>
- [428] Al Jazeera (2026, October 14). Ukrainian ex-officer goes on trial in Germany over Nord Stream blasts (title as shown; read through a tool summary). <https://www.aljazeera.com/news/2026/10/14/ukrainian-ex-officer-goes-on-trial-in-germany-over-nord-stream-blasts>

- [429] Sainz, S. (2026, June 18). Discrimination and treatment of population groups in Europe: ethnic minorities, migrants and asylum seekers (Section 11.3, revised governance roadmap: positive conditionality and assisted governance transition pacts, both proposals). ISDO Economics and Social Affairs Department. <https://isdo.ch>
- [430] United Nations Security Council (2014, March 1). Verbatim record of the 7124th meeting, S/PV.7124, on the letter of 28 February 2014 from the Permanent Representative of Ukraine (S/2014/136); read through a tool summary of a mirror page on October 22, 2026. <https://un-project.org/meeting/S-PV-7124/>
- [431] May, T. (2018, March 12). PM Commons statement on Salisbury incident: 12 March 2018 (oral statement to Parliament by the Prime Minister; read through a tool summary on October 22, 2026). UK Government. <https://www.gov.uk/government/speeches/pm-commons-statement-on-salisbury-incident-12-march-2018>
- [432] Federal Bureau of Investigation, Baltimore Field Office (2024, March 26; update of March 27). FBI Baltimore statement on Francis Scott Key Bridge collapse (read through a tool summary on October 22, 2026). <https://www.fbi.gov/contact-us/field-offices/baltimore/news/fbi-baltimore-statement-on-francis-scott-key-bridge-collapse>
- [433] EUvsDisinfo, East StratCom Task Force of the European External Action Service (2019, November 30). Disinfo: Finland provoked the war against Soviet Union in 1939, Soviet leadership stated (read through a tool summary on October 22, 2026). <https://euvsdisinfo.eu/report/finland-has-started-the-war-against-soviet-union-in-1939/>
- [434] UNHCR Refworld (consulted 2026, October 22). Charter of the United Nations (the page states that the Charter was signed at San Francisco on 26 June 1945; read through a tool summary). <https://www.refworld.org/legal/constinstr/un/1945/en/27654>
- [435] United Nations, Office of Legal Affairs, Codification Division (consulted 2026, October 22). Definition of Aggression, General Assembly resolution 3314 (XXIX), adopted 14 December 1974 (page of the Audiovisual Library of International Law, read through a tool summary; the text of the resolution was not read). <https://legal.un.org/avl/HA/da/da.html>
- [436] United Nations, Office of Legal Affairs, Codification Division (consulted 2026, October 22). Protocol Additional to the Geneva Conventions of 12 August 1949, and relating to the Protection of Victims of International Armed Conflicts (Protocol I), Geneva, 8 June 1977 (page of the Audiovisual Library of International Law; read through a tool summary). <https://legal.un.org/avl/ha/pagc/pagc.html>
- [437] National Security Archive, George Washington University (consulted 2026, October 22). The secret CIA history of the Iran coup, 1953, Electronic Briefing Book no. 28 (the page states that the history was first disclosed by The New York Times on April 16 and June 18, 2000; read through a tool summary). <https://nsarchive2.gwu.edu/NSAEBB/NSAEBB28/>

Annex A. Glossary

This glossary defines 58 terms as the manual uses them, in alphabetical order. The last column names the numbered sections, or the parts of Annex D, where a term does most of its work. The authors' own constructs (Proposals A to C, attribution posture, the two-sided procedure, and the gates and symmetry rule of Annex D) have no standing in law or doctrine. Annex E repeats the 58 terms with a Spanish rendering.

Table 31. Glossary of terms used in this manual

Term	Definition as used here	Section
Access conditions	Who could inspect the scene, logs or wreckage, and who could not	7, 8
Alleged (grade)	Asserted but unproven; the grade marks an untested claim and does not mean false	1, 7
Amplifier	Person, outlet or network that spreads an operation's content or claims	5
Anchor (geolocation, chronolocation)	Verifiable feature or record that fixes where or when an item was captured; the annex asks for two	D.1
Asymmetric warfare	Conflict between actors of unequal strength; theory asks why weaker actors sometimes win	2
Attribution	Naming the actor responsible for an act, as a legal, technical or political judgment	4, 5, 8
Attribution deception	Proposal C: any deliberate effort to mislead an audience about who is responsible	3
Attribution posture	The authors' untested variable: acknowledged, gray, denied or falsified, coded from cited sources by the rule in Section 5.8	5
Black propaganda	Products that "appear to emanate from a source other than the true one"	3
Breakout Scale	Six-category scale of an operation's spread by platform span and community span	5
Call-and-outcome register	Record of each assessed call and, later, its outcome, kept to estimate error rates	D.3, D.4
Chain of custody	Dated record of who held an item of evidence and what was done to it	D.1
Challenge review	Independent attempt to refute the leading hypothesis, recorded in a memo	D.1, D.3
Contested (grade)	Strongly evidenced, but attribution, intent or scope is disputed by credible parties	1

Term	Definition as used here	Section
Countermeasures	Legal: otherwise unlawful response to a wrongful act (ILC Article 49). Policy: any defensive step	4, 8
CSIRT	Computer security incident response team; under NIS2, the body that receives incident notifications	D.2
Cyber false attribution	Planted or borrowed artifacts that make an intrusion point to another actor	3, 4
Deception	Doctrinal term for deliberate measures to mislead targeted decision-makers	3
Decision gate	Checkpoint that opens a class of statement only when its evidence conditions are met	D.3
Deniability, plausible	Running operations so a government can plausibly disclaim them (NSC 10/2, 1948)	2, 3, 9
Documented (grade)	Confirmed by admission, court findings, declassified documents or several independent investigations	1
Due diligence	Claimed State duty not to let its territory be used against other States; disputed	4
Effective control	Attribution test requiring control of the specific operation; the ICJ declined the wider overall-control test	4
Evidence tiers	Tier 1 declared doctrine, Tier 2 verified practice, Tier 3 unconfirmed narrative	1
Evidentiary copy	Hashed, unaltered copy of a source file, used for analysis and replication	D.1
False colors	A vessel's use of another State's flag as a ruse; lawful until the attack begins	3, 4
False flag (Proposal A)	Act designed so that responsibility will be ascribed to a specific other actor, for advantage	3
False flag, pretext subtype (Proposal B)	Proposal A operation whose main function is to justify a response already intended	3
False flag, registry sense	Vessel flying a flag it is not entitled to; no blame is shifted	3, 6
FIMI	Foreign information manipulation and interference; the EEAS counts it as incidents	4, 5, 8
Gray zone	"Operational space between peace and war" where coercion stays below escalation thresholds	2
Hash value	Fixed-length digest of a file's bytes; any change to the file changes the digest	D.1
Hybrid warfare	Contested label for fused modes of fighting and threshold management; no agreed definition	2
Hypothesis log	Dated record of each hypothesis, its tests and grades, and why it was kept or dropped	D.2, D.4

Term	Definition as used here	Section
ICD 203	US directive listing nine analytic tradecraft standards, signed 2 January 2015	8
Kampala amendments	2010 amendments giving the ICC jurisdiction over aggression; active since 2018	4
Liar's dividend	Gain from falsely calling a real event fake or staged (Chesney and Citron)	7
Masked identity	Borrowed or invented persona that hides the real actor	6
Metadata state	One of five conditions of a file's metadata, from original to unresolved	D.1
Misdirection (post-attack)	Officials steering an inquiry away from the real perpetrators after an attack	6, 7
NIS2	EU Directive 2022/2555; reporting clocks of 24 hours, 72 hours and one month	4, 8, 10
Non-intervention	Rule against coercing another State's free choices; persuasion alone falls short	4
Order of volatility	Sequence for collecting evidence, from the most perishable data to the most durable	D.1, D.2
Passive collection	Gathering only what the defender holds or third parties publish, with no contact with suspected systems	D.1
Perfidy	Feigning protected status to betray an adversary's confidence; banned by AP I Article 37	3, 4
Person-based sanctions	EU listings of individuals and entities for what they did, without any State finding	4, 9
PHIA yardstick	UK probability scale of the Professional Head of Intelligence Assessment, pairing words with ranges	8
Pre-bunking	Warning people about manipulation techniques before exposure, usually by short videos or games	8
Pretext	Incident, real or staged, used to justify a response already intended	3
Refuted (grade)	Contradicted by an institution or court that examined the specific claim	1, 7
Ruse of war	Lawful deception such as camouflage, decoys and misinformation; not prohibited (AP I Article 37)	3, 4
State responsibility	Law attributing conduct to a State and its consequences (ILC Articles, 2001)	4, 9
Statement register	Log of each public statement on an incident, with its wording, ceiling and approver	D.3, D.4
Stay-behind network (Gladio)	Secret NATO-linked network concealed from parliaments; existence documented, role in attacks unproven	6, 7

Term	Definition as used here	Section
Symmetry rule	A claim of staging needs the same evidence as a claim that a named actor did it	D.3
T/C/I tests	Technical traceability, multi-source convergence, evidence of intent; each graded Y, P, N, U or R	1, 6, 8
Three Warfares	PLA doctrine of public-opinion, psychological and legal warfare; text not read	2
Two-sided procedure	The authors' method testing three hypotheses: named party acted, staged or disguised, no hostile act	8

Source: Authors' compilation from Sections 1 to 10 and Annex D; quoted phrases are from the sources cited there. AP I is Additional Protocol I of 1977 to the Geneva Conventions; ILC is the International Law Commission; PLA is the Chinese People's Liberation Army.

Annex B. Case Register

This register lists the incidents, cases and claims that Sections 3 to 9 grade, under the identifiers used there, with two boundary cases (A5 and A7) that only the register carries, and the judicial and quasi-judicial decisions on which the manual relies. Part B.1 holds the incidents, cases and claims. Part B.2 holds the courts, tribunals and inquiries, with a docket number where a source printed one and the words "unverified" or "not shown" where it did not.

Letters mark the groups. A is the historical anchors from 1931 to 1964, B is Cold War Europe, C is Russia and its neighbors, D is masked identity in cyberspace, E is poisoning, chemical weapons and sabotage, and F is maritime incidents and lawful ruses. I numbers the 38 contemporary incidents graded on the three attribution tests, in the order of their start dates, with an incident known only by year or month placed after the dated incidents of that period. N numbers the 15 claims of Section 7, with N13 split into N13a and N13b. A grade attaches to one characterization of an event and not to the event as a whole, so some events carry two or three rows, marked (i), (ii) and (iii) after the identifier.

Every row carries a tier, a status and a grade, or in the incident tables the three test grades. Status is Executed where the sources show that the operation took place as characterized, Planned only where they show a plan that was never carried out, and Alleged where the characterization is asserted and not established. Annex C defines tiers, grades and tests. Table 32 defines the type labels. Identifiers, tiers and grades match those printed in Sections 6 and 7.

Table 32. Type labels used in the register

Type	Meaning
False flag	Act designed so that responsibility falls on a named other actor
Pretext	Incident, real or staged, used to justify a response already intended
Misdirection	Officials steer an inquiry away from the real perpetrators after an attack
Deniable operation	Act by a State or its proxy that the sponsor disowns or leaves unclaimed
Masked identity	Persona or borrowed tools that hide the real operator
Plan	Proposal that was never carried out
Ruse of war	Lawful deception in armed conflict, short of perfidy
Registry sense	Vessel flying a flag it is not entitled to; no blame is shifted
Concealed structure	Secret organization kept from parliaments and oversight bodies

Type	Meaning
Attribution finding	Row grades an institution's finding on who used a weapon
Overt claim	Actor's own claim of an attack, with no deception about authorship
Alleged false flag	Claim that a named party staged an event to blame another
Counter-narrative	Competing account of who acted, offered to deflect or reverse blame
Narrative weapon	Staging claim used to discredit an opponent or explain opinion as enemy influence
Conspiracy claim	Claim that an attack, disaster or health crisis was staged, with no design shown

Source: Authors' typology, built from Section 3 and the authors' case typology. The incident tables carry no type, because the three tests grade the attribution and not the act.

B.1 Incidents graded on the three tests

The 38 contemporary incidents of B.1.1 to B.1.3 run from 2018 to 2026 and are the register graded on the three tests of Section 1: technical traceability (T), multi-source convergence (C) and evidence of intent (I), each graded Y, P, N, U or R. A test is met only when it is graded Y. On that rule 20 incidents meet no test, 12 meet one, five meet two and one meets all three. If a partly met test also counts, four meet none and 17 meet all three (Table 50, Annex C.1.2). Crimea (2014), TV5Monde (2015) and the DNC intrusion (2016) sit in tables graded by tier and grade only. The tables from B.1.4 to B.1.11 hold the historical anchors, the later cases and the claims of Section 7, which are graded by tier and grade and not by the three tests. Grades are the authors' judgment from the sources read, and every status runs to October 11, 2026, except where an update note of October 22, 2026 says otherwise (Section 6.5 and the source notes of Tables 44 and 46).

B.1.1 Contemporary incidents, 2018 to 2024 (I01 to I18)

Table 33. Contemporary incidents, 2018 to 2024, with the three attribution tests

ID	Incident, year	Attributed to, and by whom	Tier	T/C/I	Status, 11 October 2026
I01	Salisbury Novichok poisonings, 2018	GRU officers, approved by the President; UK public inquiry, 4 December 2025, via summaries	2	R/P/P	GRU officers sanctioned; Russia rejects; no trial reported
I02	Nord Stream explosions, 26 September 2022	Ukraine-linked team; state direction alleged; German Federal Prosecutor, 2 July 2026	2	P/P/U	Indicted 2 July 2026; trial due 14 October 2026

ID	Incident, year	Attributed to, and by whom	Tier	T/C/I	Status, 11 October 2026
I03	Catalangate spyware, 2022	Spanish government, alleged; Citizen Lab, 18 April 2022, via a tertiary source	2	P/R/P	Government response not retrieved
I04	912 Group personas, complaint of 17 April 2023	PRC Ministry of Public Security; US Department of Justice	2	P/P/Y	34 defendants at large
I05	North Korean remote IT-worker scheme, 2023 to 2026	DPRK state; DOJ, OFAC and an 11-country alert	2	Y/Y/Y	Ongoing; Spain named once without designee
I06	Volt Typhoon, advisory of 7 February 2024	PRC state-sponsored actors; CISA, NSA, FBI and partners	2	Y/Y/P	A Chinese body calls it a US fabrication
I07	Leyton arson, UK, 20 March 2024	Wagner-linked network; UK court and police	2	Y/P/Y	Convicted; 7 to 17 years in custody
I08	Spamouflage, 2024	Law-enforcement-linked individuals; OpenAI, Meta and the FBI	2	P/Y/P	Little authentic reach
I09	Marywilka fire, Warsaw, 12 May 2024	Russian services; Polish ministers, 11 May 2025	2	P/P/P	Trial outcome not retrieved
I10	NoName057(16) arrests, Spain, 20 July 2024	Hacktivists; CCN-CERT calls state backing "probable"	2 / 3	P/P/N	Investigation continues
I11	Incendiary parcels, July 2024	Russian intelligence; Lithuanian prosecutors and Eurojust	2	P/Y/P	Charged; others on trial in Poland
I12	Hack-and-leak against a US campaign, indicted 27 September 2024	IRGC; US Department of Justice indictment	2	P/Y/Y	Charges pending
I13	Romanian first round annulled, 6 December 2024	Russia alleged; party-paid campaign also alleged; Romanian Constitutional Court	2 / 3	P/P/P	Legal assessment disputed (Section 8)
I14	Eagle S cable damage, 25 December 2024	None; crew prosecuted; Finnish courts	2	P/N/N	Appeal court affirmed jurisdiction, 27 August 2026; appeal deadline 26 October
I15	Attempted arsons in four countries, 2024	GRU-linked group in Russia; Lithuanian prosecutors	2	P/P/Y	Six charged, 16 January 2026

ID	Incident, year	Attributed to, and by whom	Tier	T/C/I	Status, 11 October 2026
I16	Munich sabotage-planning case, 2024 to 2025	Russian intelligence; Munich court	2	P/P/Y	Convicted; six years
I17	Air-traffic-control cyberattack and Storm-1516 content, 2024	GRU (APT28) and Storm-1516; German Foreign Ministry, 12 December 2025	2	R/Y/P	Ambassador summoned
I18	Zeraati stabbing, London, 2024	Iranian state; Old Bailey, 3 July 2026	2	Y/P/Y	Two men convicted; 12 and 8 years

Source: Authors' compilation from [91, 135, 186, 187, 188, 192, 194, 196, 197, 198, 204, 211, 212, 213, 229, 231, 232, 233, 295, 301, 325, 326, 327, 328, 329, 330]. Grades are the authors' judgment, and no source publishes scores. T/C/I means technical traceability, convergence and intent; Y met, P partly, N not met, U not public, R not retrieved by this review (the evidence may be public). Rows I05, I06, I07, I09, I10, I15 and I18 also appear in Table 18. GRU is Russia's military intelligence service; IRGC is Iran's Revolutionary Guard Corps; PRC is the People's Republic of China; DPRK is North Korea; OFAC is the US Treasury's Office of Foreign Assets Control; CISA is the US Cybersecurity and Infrastructure Security Agency. The Salisbury inquiry report was read only through summaries.

B.1.2 Contemporary incidents, 2025 (I19 to I33)

Table 34. Contemporary incidents, 2025, with the three attribution tests

ID	Incident, year	Attributed to, and by whom	Tier	T/C/I	Status, 11 October 2026
I19	Hong Tai 58 cable damage, Taiwan, 25 February 2025	Chinese crew; PRC direction not found; Taiwanese court	2 / 3	P/N/N	Captain jailed three years
I20	i-Soon and APT27 intrusions, 5 March 2025	Ministry of Public Security and Ministry of State Security, via contractors; US Department of Justice	2	P/P/Y	Defendants at large
I21	Iberian blackout, 28 April 2025	No actor and no cyber cause; government, Audiencia Nacional, ENTSO-E	2	N/N/N	Provisional dismissal, 12 January 2026
I22	Operation Eastwood, 16 July 2025	NoName057(16); Europol	2	Y/P/N	Resumption not documented
I23	Authentic Antics malware, 18 July 2025	GRU Unit 26165; UK National Cyber Security Centre and Foreign Office	2	P/Y/P	Sanctions imposed
I24	Salt Typhoon intrusions, advisory of 27 August 2025	PRC-linked firms; 13-state advisory, Spain included	2	Y/Y/P	80 or more countries affected, per FBI

ID	Incident, year	Attributed to, and by whom	Tier	T/C/I	Status, 11 October 2026
I25	Drones over Poland, 10 September 2025	Russia, "deliberately"; Polish government. A Russian analyst cites a Ukrainian "false flag"	2 / 3	N/P/P	NATO Article 4 invoked; no forensics published
I26	Copenhagen airport closure, 22 September 2025	No actor; Danish police	2	N/N/N	Closed 25 June 2026; drones not confirmed
I27	Brussels airport drones, 4 November 2025	"Most likely Russia"; anonymous sources and a minister	3	N/N/P	Moscow denies
I28	Polish rail sabotage, 15 November 2025	Two Ukrainians working for Russia; Prime Minister Tusk and prosecutors	2	P/N/P	One detained; two at large
I29	GTG-1002 AI-run espionage report, November 2025	Chinese state-sponsored actor; vendor self-report	3	U/N/U	No indicators published; outside check of the grade pending (Annex C.1.4)
I30	Fitburg cable damage, 31 December 2025	None; Finnish authorities, via a tertiary source	3	P/N/N	Anchor damage reported
I31	Cyberattacks on French bodies, 2025	APT28; French Foreign Ministry, 29 April 2025	2	R/Y/P	Russian reply not retrieved
I32	Fake accounts targeting Taiwan, 2025	PRC organs; Taiwan's National Security Bureau	3	R/P/R	Evidence not described in the report relayed
I33	Baltic GNSS jamming, 2025 to 2026	Russian electronic warfare; Lithuanian regulator	2	P/Y/N	Pilot reports continue

Source: Authors' compilation from [108, 134, 148, 149, 189, 190, 191, 193, 201, 202, 203, 215, 216, 217, 218, 219, 277, 278, 279, 280, 289, 331, 332, 333, 334, 335, 336]. Rows I25, I26 and I32 also appear in Table 18. Grade R marks a source the review did not retrieve (I31, I32). I29 rests on a vendor's report of its own product and is criticized for publishing no indicators of compromise. I30 rests on a Wikipedia page relaying Finnish authorities. I32 comes from a party to the dispute, in a relayed account. FBI is the US Federal Bureau of Investigation; ENTSO-E is the European network of transmission system operators for electricity.

B.1.3 Contemporary incidents, 2026 (I34 to I38)

Table 35. Contemporary incidents, 2026, with the three attribution tests

ID	Incident, year	Attributed to, and by whom	Tier	T/C/I	Status, 11 October 2026
I34	Storm-1516 look-alike press domains, February 2026	Storm-1516, Russia; VIGINUM	2	Y/P/P	Documented before municipal elections
I35	Hormuz closure and ship-attack claims, from February 2026	Iran and the US or Israel, each blaming the other; parties, Crisis Group	3	N/N/N	October status not retrieved
I36	Liège synagogue blast and Antwerp arson, 9 March 2026	Tehran; Belgian CUTA, 24 September 2026	2	U/P/P	Suspect held in Turkey
I37	Handala persona operations, March 2026	MOIS; DOJ domain seizure, Dataminr	2 / 3	P/P/P	Persona active after seizure
I38	HAYI and IMCR attacks, from March 2026	IRGC Quds Force, "almost certainly"; UK minister, 13 July 2026	2	U/P/P	IMCR banned; Iran gave no immediate comment

Source: Authors' compilation from [195, 206, 207, 208, 209, 210, 296, 337, 338]. Row I38 also appears in Table 18. The Storm-1516 row rests on a trade-press report of VIGINUM's findings with no date recorded, and the VIGINUM analysis cited alongside it has data only to 5 March 2025. The primary VIGINUM text on the look-alike domains was not located. MOIS is Iran's Ministry of Intelligence and Security; VIGINUM is France's agency against foreign digital interference; CUTA is Belgium's threat-analysis coordination body.

B.1.4 Historical anchors, 1931 to 1964 (A1 to A9)

Group A holds nine anchors. Seven of them appear in Table 16. The CIA's Iran operation of 1953 (A5) and the Protocol of Sèvres (A7) are boundary cases that only this register carries.

Table 36. Historical anchors, 1931 to 1964: full register

ID	Case, year	Attributed to, and by whom	Type	Status	Grade; tier
A1	Mukden Incident, 1931	Kwantung Army officers; Tokyo Tribunal judgment	False flag (pretext)	Executed	Documented; 2
A2(i)	Reichstag fire, 1933: Nazi authorship	Nazi leadership; Hett and others, against the lone-arsonist school	False flag	Alleged	Alleged; 3
A2(ii)	Reichstag fire, 1933: exploitation	Nazi government; Decree of 28 February 1933	Pretext	Executed	Documented; 2

ID	Case, year	Attributed to, and by whom	Type	Status	Grade; tier
A3	Gleiwitz incident, 1939	SD team under Naujocks; Abwehr officer Lahousen at Nuremberg	False flag (pretext)	Executed	Documented; 2
A4	Shelling of Mainila, 1939	Soviet forces; Finnish record, Soviet archives, Khrushchev's memoir	False flag (pretext)	Executed	Documented; 2
A5	Operation TPAJAX, Iran, 1953	CIA; CIA history as summarized by Gasiorowski	False flag or pretext	Alleged	Alleged; 3
A6	Lavon Affair, 1954	Israeli military intelligence (Unit 131); Egyptian trial, Israeli inquiries	False flag	Executed	Documented; 2
A7	Protocol of Sèvres, 1956	France, Britain and Israel; Shlaim and participant memoirs	Pretext (collusive)	Executed	Documented; 2
A8	Operation Northwoods, 1962	Joint Chiefs of Staff memorandum; National Security Archive release, 2001	Plan	Planned only	Documented; 1
A9(i)	Gulf of Tonkin, 1964: misreported attack	US government; NSA historian's study, McNamara's 2003 acknowledgment	Pretext	Executed	Documented; 2
A9(ii)	Gulf of Tonkin, 1964: fabrication	US government, alleged; NSA study found no "smoking gun"	False flag	Alleged	Alleged; 3

Source: Authors' compilation from [54, 55, 163, 164, 165, 166, 167, 168, 169, 170, 171, 172, 173, 339, 340, 341, 342, 343, 344, 433, 437]. Mainila rests on secondary sources, among them an EU disinformation report that says Russia later admitted the staging and names no document or date [433], and the research located no formal admission. Lavon rests on press accounts and an essay [169, 170, 343]. A5 is provisional: the sham bombing is documented, but the summaries read do not say that anyone was meant to blame communists. The Wilber history, the CIA's own account of the coup, was first disclosed by The New York Times on 16 April and 18 June 2000 [437]. Neither the history nor the Times report was read at source, and A5 stays Alleged until the history is. If it shows that the bombing was staged to blame communists, A5 meets the Tier 2 test for design and joins the documented cases. Hitler's address of 22 August 1939 is a Tier 1 element of A3. The Reichstag rows rest largely on a Wikipedia compilation and on a press report of Hett's argument.

B.1.5 Cold War Europe: Italy and Turkey (B1 to B5)

Table 37. Cold War Europe, 1955 to 1980

ID	Case, year	Attributed to, and by whom	Type	Status	Grade; tier
B1(i)	Piazza Fontana, Milan, 1969: attack and misdirection	Far-right Venetian cell; Italian courts, 1987 and 2005	Misdirection	Executed	Contested; 2
B1(ii)	Piazza Fontana: state or NATO direction	Italian state or US and NATO services; authors and a 2000 report	False flag	Alleged	Alleged; 3
B2	Peteano, 1972	Ordine Nuovo cell; Venice assizes, 1987	Misdirection	Executed	Documented; 2
B3(i)	Gladio network, disclosed 1990	SIFAR and CIA; Prime Minister's disclosure, 24 October 1990	Concealed structure	Executed	Documented; 2
B3(ii)	Gladio: role in attacks	Gladio or NATO; some authors	False flag	Alleged	Alleged; 3
B4(i)	Bologna station, 1980: perpetrators and misdirection	Neo-fascist militants; Court of Cassation, 1995 and later	Misdirection	Executed	Documented; 2
B4(ii)	Bologna: financing and coordination chain	Gelli and P2; appeal court reasoning, reported 2025	Misdirection	Executed	Contested; 2
B5	Istanbul pogrom, 1955	Tactical Mobilization Group, alleged; Greek court, Turkish trials, a general	False flag (pretext)	Executed	Contested; 2

Source: Authors' compilation from [174, 175, 176, 177, 178, 179, 180, 181, 182, 183, 184, 185, 345, 346]. Court findings are known only through press and legal summaries, and docket numbers are unverified. B1(i) is Contested because two secondary accounts of the 2005 Cassation judgment conflict. B1(ii), B3(ii) and claim N8 are one proposition, graded Alleged in each place. B5 rests on Wikipedia and a Turkish news site relaying scholarship and a general's statements.

B.1.6 Russia and its neighbors (C1 to C7)

Table 38. Russia and its neighbors, 1999 to 2024

ID	Case, year	Attributed to, and by whom	Type	Status	Grade; tier
C1	Russian apartment bombings, 1999	FSB, alleged by journalists and authors; no court or commission finding	False flag	Alleged	Alleged; 3
C2(i)	Russo-Georgian war, 2008: origin of hostilities	Provocation on both sides; EU fact-finding mission, 2009	Pretext	Executed	Contested; 2
C2(ii)	Russo-Georgian war: Russia engineered the Georgian attack	Russia; Cornell and Felgenhauer, as relayed	False flag	Alleged	Alleged; 3
C3	Crimea, "little green men", 2014	Russian forces; Putin's statement of 17 April 2014	Deniable operation	Executed	Documented; 2
C4(i)	MH17 shootdown, 2014	Russian Buk crew and separatists; Dutch court and ECtHR	Deniable operation	Executed	Documented; 2
C4(ii)	MH17: Ukraine-blame accounts	Ukraine, per Russian officials and media; rejected by investigators	Counter-narrative	Alleged	Refuted; 3
C5(i)	Pre-invasion "provocations", February 2022	Separatist leaders and Russian media; Bellingcat metadata analysis	False flag (pretext)	Executed	Contested; 2 (limited)
C5(ii)	US warning of a staged Russian pretext, 2022	Russia; US officials, no evidence published	Alleged false flag	Alleged	Alleged; 3

ID	Case, year	Attributed to, and by whom	Type	Status	Grade; tier
C6	"Dirty bomb" allegation, 2022 (claim N2)	Ukraine, per Russia; IAEA inspections found no sign	Alleged false flag	Alleged	Refuted at those sites; 3
C7	Crocus City Hall attack, 2024	Ukraine, per Russian investigators; claimed by the Islamic State	Counter-narrative	Alleged	Alleged; 3

Source: Authors' compilation from [80, 154, 225, 226, 227, 228, 237, 286, 347, 348, 349, 350, 351, 352, 353, 354, 355, 356, 357]. C1 rests on journalists' and authors' work and on a partisan source, the minority staff of a Senate committee. C5(i) is limited because the central finding rests on two open-source analysts. C7 is Alleged on the evidence of Russian investigators alone.

B.1.7 Masked identity in cyberspace (D1 to D5)

Table 39. Masked or borrowed identity in five cyber cases: full register

ID	Case, year	Attributed to, and by whom	Type	Status	Grade; tier
D1	TV5Monde sabotage, 2015	GRU through APT28; French Foreign Ministry, 29 April 2025	Masked identity	Executed	Contested; 2
D2	DNC intrusion, DCLeaks and Guccifer 2.0, 2016	12 GRU officers; US grand jury, 13 July 2018	Masked identity	Executed	Contested; 2
D3	Olympic Destroyer, 2018	GRU Unit 74455; UK and US, 19 October 2020	Masked identity	Executed	Contested; 2
D4	Turla use of OilRig tools, 2019	Turla; NSA and UK NCSC, 21 October 2019	Masked identity	Executed	Documented; 2
D5	CyberAv3ngers persona, from 2023	IRGC-affiliated actors; CISA, 1 December 2023; US Treasury	Masked identity	Executed	Contested; 2

Source: Authors' compilation from [56, 60, 61, 62, 144, 219, 220, 221, 222, 223, 224]. None of these attributions has been tested at trial, the accused individuals are at large, and the accused States deny or ignore the findings.

B.1.8 Poisoning, chemical weapons and sabotage (E1 to E4)

Table 40. Poisoning, chemical weapons and pipeline sabotage, 2013 to 2026

ID	Case, year	Attributed to, and by whom	Type	Status	Grade; tier
E1(i)	Salisbury and Amesbury poisonings, 2018	GRU officers, approved by the President; UK prosecutors, EU Council, public inquiry	Deniable operation	Executed	Contested; 2
E1(ii)	Salisbury: staged by Britain (claim N6)	Britain, per Russian officials; contradicted by OPCW and EU listing	Alleged false flag	Alleged	Refuted; 3
E2(i)	Ghouta sarin attack, 2013	Syrian government forces, per governments' assessments; UN mission did not attribute	Attribution finding	Executed	Contested; 2
E2(ii)	Ghouta: rebel provocation or foreign staging	Rebels or a foreign service, per Putin and Lavrov, 2013	Alleged false flag	Alleged	Alleged; 3
E3(i)	Khan Shaykhun sarin, 2017, and Umm Hawsh mustard, 2016	Syrian government (sarin), ISIL (mustard); OPCW-UN mechanism, 26 October 2017	Attribution finding	Executed	Contested; 2
E3(ii)	Douma chlorine attack, 2018	Syrian Arab Air Forces; OPCW Investigation Team, 27 January 2023	Attribution finding	Executed	Contested; 2
E3(iii)	Douma: staged attack (Section 7)	Rebels, White Helmets or Western services, per Russian ministries	Alleged false flag	Alleged	Alleged; 3
E4(i)	Nord Stream sabotage, 2022: Ukraine-linked team	Ukraine-linked team; German prosecutors and Federal Court of Justice, 2025 to 2026	Deniable operation	Executed	Contested; 2

ID	Case, year	Attributed to, and by whom	Type	Status	Grade; tier
E4(ii)	Nord Stream: Ukrainian state direction	Ukrainian state bodies, as indicted; Ukraine denies	Deniable operation	Alleged	Alleged; 3
E4(iii)	Nord Stream: United States with Norway or Britain	United States, Norway or Britain; Hersh, 2023; Russian officials	Counter-narrative	Alleged	Alleged; 3

Source: Authors' compilation from [229, 230, 231, 232, 233, 243, 261, 262, 263, 264, 295, 358, 359, 360]. The Salisbury inquiry report was read only through summaries. The OPCW-UN mechanism's mandate lapsed on 17 November 2017 after a Russian veto, and Russia's reply to the Investigation Team's report was not retrieved. The Nord Stream trial was due to open on 14 October 2026 and opened that day (Section 6.5).

B.1.9 Maritime incidents of 2019 and 2026 (F1 and F2)

Table 41. Tanker and ship incidents in the Gulf region, 2019 and 2026

ID	Case, year	Attributed to, and by whom	Type	Status	Grade; tier
F1(i)	Fujairah tanker sabotage, 12 May 2019	"Most likely a state actor"; UAE-led investigation, 6 June 2019	Deniable operation	Executed	Contested; 2 (limited)
F1(ii)	Gulf of Oman tanker attacks, 13 June 2019	IRGC; UK statement of 14 June 2019 and US	Deniable operation	Alleged	Alleged; 3 (borderline)
F1(iii)	Gulf of Oman: staged to blame Iran	US, Israel or another third party, per Iranian officials	Alleged false flag	Alleged	Alleged; 3
F2a	Ship attacks claimed by the IRGC, March to October 2026	IRGC, by its own statements	Overt claim	Executed	Documented by claim; 2
F2b	Missiles toward Turkey, March 2026	Iran, per NATO interceptions and Turkish data; Iran denies and hints "false flag"	Deniable operation	Executed	Launch: Contested; 2. Hint: Alleged; 3
F2c	Missiles toward Diego Garcia, 19 to 20 March 2026	Iran, per UK and Israeli officials; Iran calls it an "Israeli false flag"	Alleged false flag	Alleged	Alleged; 3

ID	Case, year	Attributed to, and by whom	Type	Status	Grade; tier
F2d	Kuwaiti tanker Al-Salmi drone strike, 30 March 2026	Iran, per Kuwait Petroleum Corporation alone; no claim	Deniable operation	Alleged	Alleged; 3
F2e	Misdescribed flags of blockade-related ships, 2026	CENTCOM and Palau statements differ; sources disagree on dates	Registry sense	Executed	Contested; 2
F2f	Drone attack at Damietta, about 29 July 2026	Unattributed; Egypt's Cabinet; a Houthi source denies	Deniable operation	Alleged	Alleged; 3
F2g	Tihama attack off Yemen, August 2026	Houthis, by their own claim; Egypt disowns the ship	Overt claim	Executed	Authorship: Documented by claim; 2. Cargo claim: Alleged; 3
F2h	"Two sea mines" claim, 31 August 2026	IRGC claim; CENTCOM calls it disinformation	Counter-narrative	Alleged	Alleged; 3
F2i	Flydubai flight FZ1073, 30 September 2026	Iranian link alleged; Iran cites an "Israeli false flag" in advance	Alleged false flag	Alleged	Alleged; 3

Source: Authors' compilation from [254, 255, 256, 257, 258, 259, 260, 361, 362, 363, 364, 365, 366, 367, 368]. Every 2026 row rests on press reports and on Wikipedia-relayed context, and no institutional investigation report was found for any of them. Documented by claim means that the actor itself states authorship. F1(i) is limited because the investigating States are regional adversaries of the State later accused. F2e is Contested because the sources disagree on dates and wording.

B.1.10 Lawful ruses of war (F3 to F5)

Table 42. Lawful ruses of war, 1915 to 1947

ID	Case, year	Attributed to, and by whom	Type	Status	Grade; tier
F3	HSK Kormoran and HMAS Sydney, 1941	German raider disguised as a Dutch merchant ship; Australian committee and Cole Inquiry	Ruse of war	Executed	Disguise Documented; ensign moment Contested; 2
F4	Q-ships, 1915 to 1943	British, German and American navies; secondary account	Ruse of war	Executed	Documented; 2
F5	Operation Greif and the Skorzeny trial, 1944 to 1947	German troops in US uniforms; US tribunal acquitted Skorzeny, 1947	Ruse of war	Executed	Documented; 2

Source: Authors' compilation from [59, 369, 370, 371]. F4 rests on a Wikipedia page relaying secondary works, and its figures are uncertain. The Skorzeny tribunal gave only a finding of not guilty, so the acquittal settles no general rule on enemy uniforms.

B.1.11 Claims graded in Section 7 (N1 to N14)

Table 43. Fifteen claims graded against the institutional record: full register

ID, family	Claim examined	Examined by	Type	Status	Grade; tier
N1, a	Bucha killings were staged	OHCHR; satellite imagery reported by Reuters	Narrative weapon	Alleged	Refuted; 3
N2, a	Ukraine prepared a "dirty bomb" (case C6)	IAEA inspections at three sites, at Ukraine's request	Narrative weapon	Alleged	Refuted at those sites; 3
N3, a	A HIMARS strike caused the Olenivka explosions	OHCHR; weapon and origin undetermined	Narrative weapon	Alleged	Refuted as to weapon; 3
N4, a	Kakhovka dam and Zaporizhzhia plant attacks blamed on the other side	UN commission; IAEA, which declines attribution	Narrative weapon	Alleged	Alleged; 3
N5, a and b	Each side's account of the 3 May 2023 Kremlin drone attack	No institutional finding; analysts and anonymous US officials disagree	Alleged false flag	Alleged	Alleged; 3
N6, b	Britain staged the Salisbury poisoning (case E1(ii))	OPCW, via an AP report; EU listing of GRU officers	Alleged false flag	Alleged	Refuted; 3

ID, family	Claim examined	Examined by	Type	Status	Grade; tier
N7, b	The Pahalgam attack of 22 April 2025 was a false flag	India's chargesheet only; no independent investigation located	Alleged false flag	Alleged	Alleged; 3
N8, b	Gladio, the CIA or NATO directed the Italian bombings (cases B1(ii), B3(ii))	Italian courts; Senate commission working draft	Alleged false flag	Alleged	Alleged; 3
N9, c	The 9/11 attacks and the WTC 7 collapse were staged	9/11 Commission; NIST	Conspiracy claim	Alleged	Refuted; 3
N10, c	Sandy Hook was a staged hoax	Connecticut State's Attorney; Waterbury civil court	Conspiracy claim	Alleged	Refuted; 3
N11, c	The 2024 Key Bridge collapse was an attack	FBI; NTSB	Conspiracy claim	Alleged	Refuted; 3
N12, c	Directed-energy weapons started the 2023 Lahaina fire	Maui Fire Department and ATF	Conspiracy claim	Alleged	Refuted; 3
N13a, c	COVID-19 was a staged hoax	None located in the sources read	Conspiracy claim	Alleged	Alleged; 3
N13b, c	COVID-19 was a deliberate release	ODNI; WHO Scientific Advisory Group	Conspiracy claim	Alleged	Alleged; 3
N14, d	Russian influence alone decided the 2016 US vote	ODNI; a 1,496-person panel study	Narrative weapon	Alleged	Alleged; 3

Source: Authors' compilation from [153, 174, 175, 181, 230, 234, 235, 236, 237, 238, 239, 240, 241, 242, 243, 244, 245, 246, 247, 248, 249, 250, 251, 252, 253, 372, 432]. Identifiers, families and grades are those of Table 20. Eight claims are Refuted and seven Alleged. Refuted means an institution or court examined the specific claim; Alleged says a claim is unproven and does not say it is false.

B.2 Judicial and quasi-judicial register

This part lists the courts, tribunals and inquiries on which the manual relies. Docket numbers are given as the sources printed them. "ICJ no." is the number in the case-page address of the International Court of Justice, because the pages do not print the label General List. Where the sources read printed no number the cell says "not shown" or "not located", which does not mean that none exists, and where a number came by secondary relay or could not be checked it says "unverified" or "unconfirmed". In Tables 44 to 46, 20 of 40 rows carry a case number or other identifier, and 2 of those are flagged as unconfirmed. Most national rulings are known only through press and

legal summaries, so each must be rechecked against the judgment before quotation. In the last column, "not located" means that the research found no report.

B.2.1 International courts and tribunals

Table 44. International courts and tribunals cited in the manual

Case	Court	Docket	Year	Holding relevant here	Status of execution
Nicaragua v. United States	ICJ	ICJ no. 70; I.C.J. Reports 1986 p. 14	1986	Attribution needs effective control; the US still breached non-intervention and force rules	Reparation never fixed; removed from the List, 26 September 1991
Corfu Channel (United Kingdom v. Albania)	ICJ	ICJ no. 1	1949	State must not "allow knowingly its territory to be used" against others (second-hand quotation)	GBP 844,000 ordered, 15 December 1949; payment not verified
Oil Platforms (Iran v. United States)	ICJ	ICJ no. 90	2003	United States had not shown the conditions for self-defense	Declaratory; no payment issue
Armed Activities (DRC v. Uganda)	ICJ	ICJ no. 116	2005; 2022	Support to irregular forces breached non-intervention and force rules	USD 325 million in five installments from 1 September 2022; first reported paid, rest unverified
Bosnia and Herzegovina v. Serbia and Montenegro	ICJ	ICJ no. 91; I.C.J. Reports 2007 p. 43	2007	Acts of the VRS not attributed to Serbia; duties to prevent and punish breached	Declaratory satisfaction only
Prosecutor v. Tadić	ICTY Appeals Chamber	IT-94-1-A	1999	Overall control test, used to classify a conflict as international	Final, 15 July 1999

Case	Court	Docket	Year	Holding relevant here	Status of execution
Ukraine v. Russia, provisional measures order	ICJ	ICJ no. 182, from the case-page address	2022	No evidence substantiating Russia's genocide allegation; first two measures adopted 13 to 2	Compliance not located
Ukraine and the Netherlands v. Russia	ECtHR Grand Chamber	Nos. 8019/16, 43800/14, 28525/20, 11055/22	9 July 2025	Russia responsible for its forces and separatist groups; violations in the note	Just satisfaction adjourned; Article 46 directions; compliance not evidenced
Georgia v. Russia (II)	ECtHR Grand Chamber	No. 38263/08	2021; 2023	No jurisdiction for active hostilities of 8 to 12 August 2008; violations found afterward	Award of EUR 129,827,500 due 28 July 2023; unpaid (note)
Ukraine v. Russia (re Crimea)	ECtHR Grand Chamber	Nos. 20958/14 and 38334/18	2021; 2024	Effective control from 27 February 2014; violations and an administrative practice	Commentator expects Russia not to comply; execution not located
Georgescu v. Romania	ECtHR, committee of three judges	No. 37327/24	2025	Inadmissible; the Court did not reach the merits of the annulment (note)	Final; unanimous, 6 March 2025
Judgment, Chapter V (Mukden incident)	International Military Tribunal for the Far East	Not shown	1948 (not shown on the page read)	Mukden incident "carefully planned beforehand" by Army General Staff and Kwantung Army officers	Execution of sentences not examined

Case	Court	Docket	Year	Holding relevant here	Status of execution
Testimony of Lahousen	International Military Tribunal, Nuremberg	Transcript vol. II, p. 450	1945	Second-hand account that prisoners were dressed as Poles for Gleiwitz; evidence, no finding	Not applicable
Situation in Georgia	International Criminal Court	Situation ICC-01/15	2022	Warrants of 30 June 2022 for three individuals; investigation phase concluded 16 December 2022	Three remain at large as far as the sources show

Source: Authors' compilation from [76, 77, 78, 79, 80, 81, 95, 163, 165, 227, 351, 352, 353, 373, 374, 375, 417, 418, 419, 424, 425]. The Corfu Channel quotation comes from a footnote of the Cyber Law Toolkit, and the Corfu compensation figure has no source link in the research record. The Tadić docket number and date were read in the transcript of the judgment hearing [425], and Cassese, who presided over the Tadić Appeals Chamber, is an interested commentator. VRS is the army of the Republika Srpska. Georgia v. Russia (II): the Grand Chamber awarded EUR 129,827,500 in just satisfaction on 28 April 2023, payable by 28 July 2023. The Council of Europe register, consulted on October 22, 2026, after the evidence cutoff, shows EUR 156,683,652.22 owed at 14 October 2026, which is the award plus EUR 26,856,152.22 in default interest, so it records no payment [417]. The EUR 253,018,000 award in the same register belongs to Georgia v. Russia (IV), no. 39611/18, judgment of 9 July 2024. Ukraine and the Netherlands v. Russia, 9 July 2025: patterns of violations of Articles 2, 3, 4 § 2, 5, 8, 9, 10, 11, 13 and 14 and of Articles 1 and 2 of Protocol No. 1 from 11 May 2014 to 16 September 2022; in the MH17 application (no. 28525/20), violations of Articles 2, 3 and 13; and a breach of Article 38 [80]. Georgescu v. Romania, 6 March 2025: Article 3 of Protocol No. 1 incompatible ratione materiae, because the presidency is not part of the legislature; Article 6 inapplicable, because standing for election is a political right; Article 13 unsupported by an arguable claim; Articles 10 and 11 manifestly ill-founded for want of argument. The Court did not reach the merits of the annulment [418, 419]. ICC-01/15 is the reference printed on the arrest warrant of 30 June 2022 [424].

B.2.2 National and military courts, historical cases

Table 45. National and military courts on historical cases

Case	Court	Docket	Year	Holding relevant here	Status of execution
Reichstag fire trial	Court at Leipzig (name not shown)	Not shown	1933	Four of five defendants acquitted; van der Lubbe condemned	Executed 10 January 1934; conviction annulled, January 2008
Trial of Skorzeny and others	General Military Government Court, Dachau	None shown	1947	Acquitted on 9 September 1947, with no reasons given	Final; nothing enforced

Case	Court	Docket	Year	Holding relevant here	Status of execution
Lavon Affair trial	Egyptian court (type not stated)	Not shown	1954 to 1955	Trial of 11 December 1954 to 27 January 1955 ended in two hangings	Hangings of 31 January 1955; Israel recognized members, 30 March 2005
Piazza Fontana, anarchists	Court of Cassation	Not verified	1987	Anarchists' acquittals confirmed, 27 January 1987	Final
Piazza Fontana, Venetian group	Court of Cassation	Not verified	2005	Sources conflict: Venetian group's part "not proven", or its responsibility affirmed	Freda and Ventura, finally acquitted earlier, could not be retried
Peteano	Corte d'Assise, Venice	Not verified	1987	Life terms for two Ordine Nuovo members; three officers convicted of concealment	Officers' sentences pardoned
Bologna station bombing	Court of Cassation	Not verified	1995	Fioravanti and Mambro life terms; Gelli and others convicted of false accusation	Final, 23 November 1995
Bologna station, Bellini	Court of Cassation, sixth section	Not verified	2025	Life sentence upheld; Segatel six years for misdirection	Reported 1 July 2025
Moscow apartment bombings trial	Moscow City Court	Not verified	2004	Life terms for two defendants after a closed trial; 243 killed	Appeal outcome unverified; a later report says no court convicted anyone

Source: Authors' compilation from [59, 169, 174, 176, 181, 185, 339, 343, 346, 347, 348]. The Leipzig and Lavon rows rest on secondary compilations. The two reports on the Moscow trial conflict, and the research could not resolve them.

B.2.3 National proceedings on contemporary incidents and claims, 2013 to 2026

Table 46. National proceedings on contemporary incidents and claims

Case	Court	Docket	Year	Holding or charge	Status of execution
United States v. Internet Research Agency and others	US District Court, D.D.C.	Criminal Action No. 18-cr-0032-2 (DLF)	2018	Indictment of 16 February 2018: 13 individuals and three entities	Justice Department moved to drop Concord charges, 16 March 2020; others not found
Indictment of 12 GRU officers (DNC intrusion)	US District Court, D.D.C.	No. 1:18-cr-00215 (ABJ), from the ECF header	2018	Charges of 13 July 2018, including the DCLeaks and Guccifer 2.0 personas	Defendants at large
Indictment of six Unit 74455 officers	US District Court, W.D. Pa.	Number not located in the texts read	2020	Seven counts; returned 15 October, announced 19 October 2020	Custody status not stated
United States v. Certain Domains	US District Court, E.D. Pa.	No. 24-mj-1395, Doc. 4	2024	FBI affidavit alleging a Russian influence operation; allegations only	Copy read ends at paragraph 107, with no exhibits
MH17 criminal trial	District Court of The Hague	ECLI not located; four judgments of 17 November 2022	2022	Girkin, Dubinsky and Kharchenko: life in absentia; Pulatov acquitted	Kremlin calls the rulings "null and void"
Eagle S prosecution	Helsinki District Court; Helsinki Court of Appeal	Not located	2025; 2026	District court dismissed, October 2025; appeal court affirmed jurisdiction, 27 August 2026	Appeal deadline 26 October 2026; cable owners seek about EUR 105 million

Case	Court	Docket	Year	Holding or charge	Status of execution
Nord Stream detention appeal	Federal Court of Justice, 3rd Criminal Senate	StB 60/25	2025	German jurisdiction; no functional immunity for intelligence-directed violence; headnotes only	Order of 10 December 2025; Wikipedia gives 23 January 2026
Trial of Serhii K. (Nord Stream)	Hanseatic Higher Regional Court, Hamburg	"3 St 2/26" per a regional outlet; court page not read	2026	Indictment admitted 2 September 2026; joint-plaintiff request refused 30 September	Due 14 October 2026; opened that day (update of 22 October)
Zeraati stabbing trial	Old Bailey	Not verified	2026	Two men convicted, 3 July 2026; judge found the attack served a foreign power	Sentences of 12 and 8 years
Leyton arson prosecutions	Old Bailey	Number not located	2025	Prison terms of 7 to 17 years, plus extended licence (note)	First convictions under the National Security Act 2023
Munich sabotage-planning trial	Munich court	Not shown	2025	Six years for a German-Russian defendant, 30 October 2025	Sentenced
Hong Tai 58 cable case	Taiwanese court	Not shown	2025	Captain jailed three years for the cable damage, June 2025	Sentenced; PRC direction not found

Case	Court	Docket	Year	Holding or charge	Status of execution
Iberian blackout inquiry	Audiencia Nacional	Not shown	2026	Provisional dismissal, 12 January 2026: no "mínimo indicio" of terrorist sabotage	Provisional
Crocus City Hall trial	Military court or Moscow City Court (sources conflict)	Not shown	2026	Verdict of 12 March 2026: 19 convicted, 15 to life	Investigators claim a Ukrainian role; verdict reportedly silent on the Islamic State
Sandy Hook defamation judgment	Connecticut Superior Court, Waterbury	Appeals AC 46131 to 46133; trial docket unconfirmed	2021 to 2025	About USD 1.4 billion on default; USD 150 million vacated on appeal	US Supreme Court rejected the appeal, 14 October 2025
French arrest warrant for Bashar al-Assad	Cour de cassation	No docket in source	2025	Annulled on head-of-state immunity; new warrant possible; facts not examined	Ruling of 25 July 2025
Annulment of the 2024 presidential election (Romania)	Constitutional Court of Romania	No. 32/6.12.2024; file 3771F/2024	2024	Annulled the 2024 presidential electoral process and ordered it rerun	Final; ECtHR declared the challenge inadmissible, 6 March 2025

Source: Authors' compilation from [61, 91, 128, 129, 132, 133, 144, 157, 186, 194, 221, 228, 232, 233, 248, 279, 329, 331, 357, 360, 372, 376, 377, 420, 421, 422, 423, 426, 427, 428]. Docket numbers appear as the sources printed them, and where sources conflict the row says so. The sources differ on whether the Crocus court was a military court or the Moscow City Court, and the claim that the verdict names Ukraine is the Russian Investigative Committee's own. The Nord Stream trial was due to open on 14 October 2026, three days after the evidence cutoff, and opened that day [428]; the file number comes from a regional outlet [427]. Leyton: six defendants were sentenced at the Old Bailey on 23 and 24 October 2025 to custodial terms of 17, 12, 9, 9, 8 and 7 years, each with an extended licence of 1 to 6 years [421]. Sandy Hook: the Appellate Court of Connecticut decided the appeals on 10 December 2024 (229 Conn. App. 487) [422]; a case summary reports that it upheld the trial court's findings except the CUTPA violation and vacated USD 150 million in statutory punitive damages [423]. The US Supreme Court declined review on 14 October 2025 [372]. The Romanian Constitutional Court decision is known from a legal note [420]. No ECLI was located for the four MH17 judgments [426].

B.2.4 Inquiries, commissions and fact-finding bodies

Table 47. Inquiries, commissions and fact-finding bodies cited in the manual

Case or inquiry	Body	Reference	Year	Finding relevant here	Status
Khan Shaykhun and Umm Hawsh	OPCW-UN Joint Investigative Mechanism	Report to the Security Council, 26 October 2017	2017	Syrian government responsible for sarin at Khan Shaykhun; ISIL for mustard at Umm Hawsh	Mandate lapsed 17 November 2017 after Russian vetoes
Douma Fact-Finding Mission	OPCW	S/1731/2019, 1 March 2019	2019	Reasonable grounds that chlorine was used as a weapon; no responsibility named	Followed by the Investigation Team's third report, 27 January 2023
Douma Investigation and Identification Team	OPCW	Third report, 27 January 2023	2023	Reasonable grounds that Syrian Arab Air Forces were the perpetrators	Russia's reply not retrieved
Ghouta mission	UN Mission to Investigate Alleged Uses of Chemical Weapons	Report of 16 September 2013	2013	Sarin used; the mandate excluded attribution	Resolution 2118 followed, 27 September 2013
Tagliavini Report	Independent International Fact-Finding Mission on the Conflict in Georgia	Report, Volume I, September 2009	2009	Georgian attack unjustifiable; Russian response went beyond reasonable defense	Both governments rejected parts
Dawn Sturgess Inquiry	UK statutory public inquiry	Final report, 4 December 2025	2025	GRU members carried out the attack, approved by the President; read through summaries	Russia rejects the attribution

Case or inquiry	Body	Reference	Year	Finding relevant here	Status
Lavon inquiries	Israeli inquiries (Olshan-Dori; Cohn)	Olshan-Dori, year not stated; Cohn, 1960	1960	No conclusive evidence Lavon authorized the operation; Cohn found perjury	Responsibility still disputed; transcripts released 11 May 2015
Italian Senate commission working draft	Commissione stragi	Doc. XXIII n. 64, 9 January 2001	2001	Institutional misdirection certain; the Stay Behind suspicion not validated	Internal draft; no adopted report
Gladio reports	Italian parliamentary commission	Doc. XXIII n. 32 (17 April 1991); n. 42 (4 December 1991)	1991	Hidden structure in the military service, dissolved 27 November 1990	SISMI director invoked state secrecy
Gladio resolution	European Parliament	Official Journal reference not retrieved	1990	Clandestine parallel intelligence and armed organization in several Member States	Asks governments to dismantle such networks; execution not examined

Source: Authors' compilation from [170, 175, 178, 179, 231, 262, 263, 295, 342, 345, 350, 358, 359]. The Dawn Sturgess Inquiry report was read only through summaries. The Italian Senate document is a working draft written by the commission's president and was never adopted. SISMI was Italy's military intelligence service.

Annex C. Methodology and source verification log

This annex states the rules by which the manual classified its evidence and logs how the sources on the commissioning list were checked. C.1 gives the procedure and its limits. C.2 records the result for each of the 34 entries on the commissioning list and for the seven narrative families set for review. All checks date from October 11, 2026, the evidence cutoff, and a later reader should expect statuses to have moved.

The annex has two uses. An auditor can trace a grade, a tier or a count in the body to the rule that produced it. The authors of a later edition can see which sources still await a first-hand reading before they can carry weight.

C.1 Verification procedure

C.1.1 Tiers and grades

Every statement in the manual rests on one of three tiers (Table 48). A tier names the kind of evidence behind a statement and says nothing about how likely the statement is to be true. Declared doctrine (Tier 1) shows what an actor acknowledges it can do, and it never proves that a given operation took place. Tier 2 needs primary or forensic records, peer-reviewed research or converging independent attributions. Tier 3 holds the rest and appears only as a disputed claim. The authors' own readings of the evidence carry the label *Inference* and sit at Tier 3.

Table 48. The three evidence tiers and their treatment

Tier	Kind of evidence	Treatment in the manual
1	Declared doctrine: an official text that acknowledges a capability or method	Shows what an actor says it can do; proves no operation
2	Verified practice: primary or forensic records, peer-reviewed research, converging attributions	May support a finding; the source line names how it was read
3	Unconfirmed narrative: inference, anonymous or partisan sourcing, no traceability	Shown only as a disputed claim, never as a fact

Source: Authors' evidence scheme (Section 1).

Grades attach to a stated characterization of a case and not to a whole event (Table 49). One event can therefore fill several rows of Annex B. The Reichstag fire is Alleged as to Nazi authorship (A2(i)) and Documented as to its political exploitation (A2(ii)). An event is classified as a false flag only where Tier 2 evidence of design exists. Otherwise it is alleged and stays at Tier 3 (Section 3). A claim is graded Refuted only where an institution or court examined that specific claim (Section 7). Status is a separate label:

Executed, Planned only or Alleged records whether the operation happened as characterized.

Table 49. Case grades and the evidence each requires

Grade	Applied when	Example in Annex B
Documented	Admission, court findings, declassified documents or several independent investigations	Gleiwitz incident, 1939 (A3)
Contested	Strongly evidenced, but attribution, intent or scope is disputed by credible parties	Piazza Fontana, 1969: attack and misdirection (B1(i))
Alleged	Asserted or unproven; untested, and not for that reason false	Reichstag fire, 1933: Nazi authorship (A2(i))
Refuted	Contradicted by an institution or court that examined the specific claim	Bucha killings were staged (N1)

Source: Authors' evidence scheme (Section 1); identifiers and grades as printed in Annex B.

C.1.2 Attribution tests and the counting rule

Three tests grade an attribution to a named actor: technical traceability (T), multi-source convergence (C) and evidence of intent (I). Each is graded Y (met on the public record), P (partly met), N (not met), U (the evidence is not public) or R (the review did not retrieve the source, so the evidence may be public). U describes the evidence and R describes the review. A grade is U only where a source says the evidence is not public or the published record shows none; where the review read a source through a summary or a relay, or could not retrieve a reply, the grade is R. Of the 114 grades, 25 are Y, 55 P, 23 N, 5 U and 6 R. The five U grades are the intent test of I02, the technical and intent tests of I29 and the technical tests of I36 and I38. The six R grades are the technical test of I01, I17, I31 and I32, the convergence test of I03 and the intent test of I32. The tests grade the attribution and not the physical event, so a well-documented incident can still carry weak grades when no source ties it to an actor.

The grades are the authors' judgment from the sources read, and no published scores exist. All 114 grades (38 incidents by three tests) are one reader's reading of the public record, and no second-rater check has been made. The authors plan to code the 114 cells blind against a written rubric, to report their agreement and to have an outside reader settle disagreements, and the counts below may then change. In summary counts a test is met only when it is graded Y, so P, N, U and R all count as not met. On that rule one of the 38 incidents meets all three tests, six meet at least two and 20 meet none (Table 50). If a P counts as met, 17 meet all three, 29 meet at least two and four meet none, and those four have no Y and no P. The R grades set a further bound.

Five incidents carry one, and if every R cell were met, three more incidents (I17, I31 and I32) would meet at least two tests and the group meeting none would fall from 20 to 17, while no incident would join the one that meets all three.

Table 50. Distribution of the 38 incidents by the number of tests met, with and without credit for P

Tests met	Y only	P counts as met	Incidents under the Y-only rule
All three	1	17	I05
Exactly two	5	12	I06, I07, I12, I18, I24
Exactly one	12	5	I04, I08, I11, I15, I16, I17†, I20, I22, I23, I31†, I33, I34
None	20	4	I01†, I02, I03†, I09, I10, I13, I14, I19, I21, I25 to I30, I32†, I35 to I38
Total	38	38	All incidents of Annex B.1.1 to B.1.3

Source: Authors' count from the grades printed in Annex B.1.1 to B.1.3. The Y-only counts equal those in the Key finding of Section 6 and in Table 1: one incident meets all three tests, six meet at least two (5 plus 1) and 20 meet none. A dagger marks an incident with a grade R. In the column headed P counts as met, a P is treated as a Y, and the four incidents with no Y and no P are I21, I26, I29 and I35.

C.1.3 Evidence handling

How a source was read limits what it may carry. The authors recorded the reading route for each source, and the manual lets the route set the use (Table 51). A source read through one route only, at abstract level, through a secondary account or as a lead carries no central finding, a term that Section 1.4 defines. The text states the limit beside the datum it qualifies, and the reference in the Bibliography names the route.

Table 51. Reading routes and the rule applied to each

Reading route	Example	Handling rule
Primary text read directly	Memorandum Opinion of 13 August 2018 in the Internet Research Agency case	May carry a central finding
Primary text through excerpts or a summarizing tool	NATO AJP-10.1, read on a truncated text	May support a finding; recheck quotations and pinpoints against the original file
Abstract or catalog record only	Person, Kulalic and Mayle (2024), read at abstract level	Supports the stated point only; no central finding
Secondary relay	A press report of a court finding or a legal summary	Corroboration; the relay is named in the text; no central finding
Wikipedia page	Pages on a journal or an author used for background	Pointer to primary references; never the sole source of a finding

Reading route	Example	Handling rule
Vendor or civil-society report	The GTG-1002 report, a vendor self-report (I29)	Corroboration only; graded Tier 3 where nothing else supports it
Title or web address only	References whose title was reconstructed from a web address	Reference marked as reconstructed; no finding drawn from it

Source: Authors' reading log; examples from [22, 126, 133, 217].

Conflicts followed one rule. Where two notes disagreed on a docket number, a date or a count, the manual gives the version read in a primary document. Counts that no primary document settled appear with their source and a hedge, or not at all. Docket numbers are printed exactly as read, and Annex B.2 marks those that remain unverified.

C.1.4 Limits

Several limits bind every finding in the manual. Many primary documents, treaties and doctrine among them, were read through extracted excerpts or summarizing tools, so quotations and pinpoint citations need rechecking against the original file before verbatim use. Web of Science could not be queried. An indexing status of "not confirmed" for it means only that the journal did not name it and the check could not be run.

The incident sample is selective, and the manual publishes no search protocol, date window or inclusion rule for it, so a reader cannot judge selection bias from the tables. What the tables show is this. The 38 incidents run from 2018 to 2026. In 34 of them a government, prosecutor, court or vendor named an actor, and four record none (I14, I21, I26, I30): they enter as cases of damage or speculation in which no authority attributed the act. Almost every source is a government or a prosecutor describing an adversary. Industry and civil-society material serves as corroboration. Peer-reviewed outcome studies are few and narrow. No operator-internal documents were seen except those quoted by prosecutors, and nothing from Chinese or Iranian doctrine on measuring effect. Counts from the register describe these 38 incidents and give no rate for contemporary incidents in general. The sample therefore shows how long denial lasted where it ended and says nothing about operations that were never exposed (Section 6).

Accused parties' replies were retrieved for few incidents, and the tables say where none was found. All 114 test grades are one reader's judgment until the blind second coding of Annex C.1.2 is done. Events after October 11, 2026 lie outside the record, so

every trial, listing and investigation may have moved. Gender and age disaggregation is not covered in this edition (Section 1.3).

Disclosure. The drafting used an AI model whose developer, Anthropic, is also the vendor behind incident I29, the GTG-1002 report. I29 was graded on the published report and its press coverage only, it stays at Tier 3 as a vendor self-report, and an outside reader has not yet checked its treatment.

C.2 Verification log

C.2.1 Method and outcome counts

The commissioning list held 34 sources, 17 filed under Tier 1 and 17 under Tier 2 by the list's own tier labels, and seven narrative families for Tier 3. The verification groups the 34 sources in five strands: concepts and doctrine (14 entries), case studies (5), threat actors (9), tactical outcome studies (5) and law (1). Each entry was checked for existence, authors, year, venue and findings. The manual follows the verified record and never the one-line description on the list.

Four outcomes are used. An entry is *As listed* when the record matches it, allowing a minor variant of the title or a caveat about the quality of the venue. It is a *Mismatch* when the year, number, author, venue or kind of document differs from the record. It is *Candidate found* when a likely match turned up that the entry did not name, and *Not found* when repeated searches returned no match. Table 52 gives the counts.

Indexing was checked through the public interface of the Directory of Open Access Journals (DOAJ) by ISSN and through a SCImago ISSN search, which draws on Scopus data. Five regional-journal items show no confirmed Scopus, Web of Science or DOAJ indexing and carry no central finding. The Tier 2 entry on the *American Journal of Economics and Sociology* and the narrative family on NATO and US false flags in Europe concern the same article, by Griffin and Cobb (2023).

Table 52. Outcome of the source check, by strand

Strand (entries)	As listed	Mismatch	Candidate found	Not found
Concepts and doctrine (14)	9	3	1	1
Case studies (5)	0	5	0	0
Threat actors (9)	7	2	0	0
Tactical outcome studies (5)	2	1	0	2
Law (1)	1	0	0	0
Total (34)	19	11	1	3

Source: Authors' recount from Tables 53 to 57, where each row is one entry.

Fifteen of the 34 entries (11 mismatches, one candidate match and three not found) could not be used as described, which is why the manual follows the verified record throughout. Nine of the fifteen are not cited in the manual and six are cited from the verified record, as the last column of Tables 53 to 57 shows.

C.2.2 Concepts and doctrine (14 entries)

The strand holds NATO doctrine, RAND reports, one indexed journal article and five regional-journal papers, and its evidence weight differs sharply. NATO doctrine is strong evidence of what is declared and says nothing about practice. The RAND reports and the *International Affairs* article are citable Tier 2 corroboration for concepts and assessments. The five regional papers add little that the stronger sources lack. Two of them (entries 11 and 13) take positions on the contested side of the literature reviewed in Section 2: one assigns a hostile campaign to named states without traceable evidence in its abstract, and the other calls hybrid warfare "new".

Inference. All three mismatches are errors of date or number: a superseded publication number and two years that are three and two years off. That pattern suggests the list was compiled from catalog or repost dates, so every other bibliographic detail in it was checked before it entered the Bibliography (Tier 3).

Table 53. Verification of the 14 concept and doctrine entries

No.	As listed	Verified record	Outcome	Use in this manual
1	NATO AJP-3.10, information operations, 2024	Superseded by AJP-10.1 (2023); no 2024 edition of AJP-3.10 found	Mismatch (number)	AJP-10.1 cited in Section 5
2	NATO AJP-10.3, military public affairs, 2024	Edition A Version 1, UK publication 11 October 2024; body not read	As listed	Not cited

No.	As listed	Verified record	Outcome	Use in this manual
3	NATO approach to countering information threats, 2024	Official text endorsed 18 October 2024; ABCDE framework in paragraph 19.1	As listed	ABCDE row in Section 5 models
4	Cognitive Warfare, NATO Science and Technology Organization	Chief Scientist report by Blatny and Søndergaard, 2025; entry under-specified	As listed	Not cited
5	RAND, NATO Strategic Communications concept, 2010	Reding, Weed and Ghez, TR-855/2-MOD/FR, 12 November 2010	As listed	Not cited
6	RAND, Lessons from Others, 2015	Paul and others, RR-1925/1-A, published 7 June 2018	Mismatch (year)	Not cited
7	RAND, Hybrid Warfare in the Baltics, 2019	Radin, RR-1577-AF, published 23 February 2017	Mismatch (year)	Typology used in Section 2
8	Person, Kulalic and Mayle, International Affairs, 2024	International Affairs 100(4), July 2024; read at abstract level	As listed	Dimensional proposal in Section 2 (abstract only)
9	Hybrid warfare case study, Syria and Crimea, IJSSHR, 2025	Aziz and Zulkifli, 8(6), June 2025; indexing unconfirmed; weaknesses from one-tool reading	As listed (weak venue)	Not cited
10	Asymmetrical warfare, Pakistan Social Sciences Review, 2025	Jamil, Fatima and Shahbaz, 9(1), 2025; HEC Y category; Scopus not confirmed	As listed (regional venue)	Not cited
11	Hybrid warfare in the 21st century, JDSS, 2024	Qureshi and Khalid, 5(4), 2024; names states without traceable evidence in abstract	As listed (regional venue)	Not cited
12	Indonesia total defense article, journal not given	Prasetyo and others, IJHESS 5(1), 28 August 2025; title longer than listed	As listed (title variant)	Not cited
13	Nimit Mai Review, Thailand, 2024, no title given	Likely match: Panngam, "Superpowers and Hybrid Warfare", Year 7, issue 1	Candidate found	Not cited
14	Drone warfare and South Korea's strategic tasks, Journal of Peace and Unification	No match after four rounds of English and Korean searches	Not found	Not cited

Source: Authors' verification, 11 October 2026, from [22, 23, 126, 378, 379, 380, 381, 382, 383, 384, 385, 386, 387, 388]. Entries 1, 4, 9 and 12 rest in part on a summarizing-tool reading (entry 1 on truncated text), the body of entry 2 was not read, and entry 8 was read at abstract level. Entry 14 has no source because nothing matched; related drone items found along the way are different works and are not cited.

C.2.3 Case studies (5 entries)

All five entries differ from the record in venue, title or kind of document. Entry 1 sits in *International Interactions* and not in the *Journal of Conflict Resolution*. Entry 2 is a press release. Entry 4 is a book chapter. Entry 5 was first published by PolitiFact. Entry 3 is the disputed Griffin and Cobb article, which also supplies narrative family (g) in Table 58, and the manual cites it for no finding (Section 7).

Table 54. Verification of the five case-study entries

No.	As listed	Verified record	Outcome	Use in this manual
1	Levin, "How Low Can You Go?", <i>Journal of Conflict Resolution</i> , 2025	Levin, <i>International Interactions</i> 51(3), online 10 April 2025; abstract only	Mismatch (venue)	Sections 3 and 9, abstract level
2	Turns, Cranfield University, 2023	Press release of 31 March 2023, republished from <i>The Conversation</i>	Mismatch (kind of document)	Not cited
3	"False-Flag Operations", <i>American Journal of Economics and Sociology</i> , 2023	Griffin and Cobb, "NATO/US false-flag attacks in Europe", 82(5), November 2023	Mismatch (title)	Section 7, cited for no finding
4	Gambetta, "The Red Brigade Signature", Oxford Academic, 2024	Chapter in <i>Fight, Flight, Mimic</i> , Oxford University Press, 2024; abstract partly visible	Mismatch (book chapter)	Not cited
5	McCarthy and Jacobson, Poynter, 2022	PolitiFact original of 8 February 2022; Poynter republished it on 9 February	Mismatch (original publisher)	Section 6, corroboration only

Source: Authors' verification, 11 October 2026, from [63, 173, 272, 273, 389, 390, 391]. Publisher pages at Taylor & Francis, Wiley, doi.org, Semantic Scholar and OpenAlex returned errors and were not retried, so entry 1 rests on the author's own page and Crossref metadata. Sample sizes, the chapter's findings and the full text of entry 3 were not seen. Wikipedia lists a co-author of entry 3 as editor-in-chief of the journal (year unchecked).

C.2.4 Threat actors (9 entries)

All nine entries exist. Two differ from the record in the kind of document: the 2013 study listed as a Department of Defense publication is a contractor study, and the 2025 "report" is a journal article adapted from entry 3. The China-side items are doctrine or forecasts, and the Russia-side items are doctrine or weakly evidenced

analysis. None documents a specific verified operation, so none supports an attribution in Section 6. Entries 3 and 4 share a lead author and count once as evidence.

Inference. Entries 1, 2 and 6 are declared doctrine or its analysis (Tier 1) and say nothing about practice. Entries 3 and 4 are forecasts by one author. Entries 5, 8 and 9 rest on thin evidence or on authors close to a party in the war, and they carry no central finding. Entry 7 is an insider's advocacy text that shows a bloc's declared stance and no practice (Tier 3).

Table 55. Verification of the nine threat-actor entries

No.	As listed	Verified record	Outcome	Use in this manual
1	PLA Political Work Regulations and the Three Warfares	Verified indirectly; approval 2003 per Halper study, 2005 per Kania; text not read	As listed (indirect)	Section 2 uses 2005
2	China's Three Warfares, US Department of Defense, 2013	Contractor study for the Office of Net Assessment, May 2013, released in litigation	Mismatch (kind of document)	Not cited
3	Gershaneck, "China's Second Battlefield", JAMS, 2024	JAMS 15(2), 2024; double anonymous peer review; no verified operation documented	As listed	Not cited
4	Political Warfare against Intervention Forces, Department of Defense, 2025	Journal of Indo-Pacific Affairs article by Gershaneck and Chan, adapted from entry 3	Mismatch (kind of document)	Not cited
5	"Grey Zone Escalation", Insignia, 2025	Sinaga, 12(2), November 2025 issue, online 7 February 2026; names no incident	As listed (thin evidence)	Not cited
6	Eggen, "A Strategy for the Weak", Defence Studies, 2025	Defence Studies 26(2), online 19 September 2025; abstract read, full text not	As listed	Not cited
7	CSTO information and psychological security policy, RUDN Journal, 2025	Vykhodets, 25(1), 2025; insider advocacy showing a bloc's declared stance	As listed	Not cited

No.	As listed	Verified record	Outcome	Use in this manual
8	Psychological operations in the Ukraine war, Land Forces Academy Review, 2024	Stancu and Căzănaru, 29(3), September 2024; reports no case data	As listed	Not cited
9	Russia's strategic information operations against Ukraine, 2026	Hurkovskyi and others, 2(66), 30 June 2026; subtitle missing from entry	As listed (title variant)	Not cited

Source: Authors' verification, 11 October 2026, from [27, 392, 393, 394, 395, 396, 397, 398, 399]. The Chinese text of the Regulations was not retrieved. The full texts of entries 6 and 9 were not read. Indexing evidence is thin for entries 3, 5 and 8: DOAJ returned no record for entry 5 and a title search found none for entry 8. The title of entry 5 is kept as printed by the journal.

C.2.5 Tactical outcome studies (5 entries)

Three entries were verified and two were not found. The three verified items study what moves public support for war: duration, expected success and civilian harm. None studies propaganda or influence operations. The set therefore shows competing determinants of support and cannot show that influence operations have no effect. Section 5 counts the five entries and draws no finding from them.

The sample sizes of entries 1 and 3 are unknown. The two missing titles may exist under other names and should be re-supplied with DOIs before use.

Table 56. Verification of the five tactical outcome-study entries

No.	As listed	Verified record	Outcome	Use in this manual
1	Sanaei, "Time is of the essence", Journal of Peace Research, 2019	JPR 56(6): 783-796; randomized survey experiment; sample not seen	As listed	Counted in Section 5; no finding drawn
2	"Limiting Civilian Casualties", Daedalus, 2017, by Kocher and others	Felter and Shapiro, Winter 2017; analytical essay without an experiment	Mismatch (author, design)	Counted in Section 5; no finding drawn
3	"Success Matters", Belfer Center, Feaver, Gelpi and Reifler	International Security, Winter 2005/06; abstract level; sample not seen	As listed	Counted in Section 5; no finding drawn
4	"Sacrificing Civilians to Save Soldiers", Terrorism and Political Violence, 2025	No match after six or more searches; nearest record has another title	Not found	Not cited
5	"Academic Considerations on the Electoral Costs of Wars of Attrition", IEEE, 2026	No match after four or more searches	Not found	Not cited

Source: Authors' verification, 11 October 2026, from [400, 401, 402]. For entry 4 the nearest topical record is a 2025 article in another journal under a different title; it was not treated as the listed work. The volume and pages of entry 3 were inferred from a file name and are unconfirmed.

C.2.6 Law strand: the sonic-boom article (1 entry)

The list also named an article on sonic booms as a psychological-warfare tool. The law strand verified it separately, and it is the one entry of the fifth strand in Table 52. It counts as *As listed*, because the year range given (2024-2025) includes the year of the record and the list shortened the title with an ellipsis. The record is Abi Chedid (2025), *International Review of Law* 14(2), published on 3 October 2025 by Qatar University Press [100]. The description supplied differed from it in three details, none of which changes the identity of the work. The year is 2025 and not 2024-2025. The article concerns one tool, sonic booms from military aircraft, rather than psychological-warfare tools in general. And the full title continues "Under International and Humanitarian Law". The abstract and the conclusion of the article disagree: the abstract finds no explicit rule, and the conclusion calls the practice "a prohibited method of warfare". Section 4 uses the abstract's finding only, as an example of the regulatory gap, and does not rely on the conclusion.

Table 57. Verification of the law-strand entry

No.	As listed	Verified record	Outcome	Use in this manual
1	Analytical Study on the Legality of Use of Sonic Boom as a "Psywar" Tool, 2024-2025	Abi Chedid, International Review of Law 14(2), 3 October 2025; fuller title	As listed	Section 4, example of the regulatory gap only

Source: Authors' verification, 11 October 2026, from [100]. The article was read once through a summarizing tool and its quotations were not rechecked.

The journal is listed in DOAJ with double anonymous peer review. Its own indexing page names Crossref, DOAJ and several legal databases but not Scopus or Web of Science, and SCImago returned no result for either ISSN. The article was read once through a summarizing tool, so quotations and page numbers need rechecking. Section 4 cites it only as a peer-reviewed example of the regulatory gap for non-kinetic psychological effects, and never as authority on false flags, ruses or State practice.

C.2.7 Narrative families

Seven families of Tier 3 narratives were set for review. Four are defined in full in Section 7 as families (a) to (d). Two are defined by topic only: Chinese psychological operations in Taiwan, the South China Sea or elections (e), and Russian information operations outside Ukraine (f). The seventh (g), the thesis that NATO and the United States staged attacks in Europe, is reconstructed here from the subject of the disputed journal article (Section 7).

Table 58 states what the record shows for each family and whether the family, as supplied, fits it. *Holds* means the framing matches the record. *Adjusted* means the framing needed correction. *Open* means the record is too thin to judge. Where the cells name claims (N1 to N14), the grades are those of Table 20.

Table 58. Verification of the seven narrative families

Family as supplied	Verified record	Fit with the record	Use in this manual
(a) Reciprocal false-flag accusations in the Ukraine war	Institutions with access settled N1 to N3 in part; N4 and N5 stay open	Holds	Section 7, claims N1 to N5; Annex B
(b) Claims that the attacked state organized the incident	Claim is that the victim state staged it; N6 Refuted, N5, N7, N8 Alleged	Holds	Section 7, claims N5 to N8; Annex B

Family as supplied	Verified record	Fit with the record	Use in this manual
(c) Attacks, disasters and health crises presented as staged	N9 to N12 refuted by investigators; N13a and N13b stay Alleged; the origin question is accident versus spillover	Adjusted (N13a and N13b touch an origin dispute)	Section 7, claims N9 to N13b
(d) Psychological warfare alone explains shifts in opinion	ODNI made no impact assessment; a 1,496-person panel found no meaningful link to vote choice	Adjusted ("alone" and "no effect" both unsupported)	Section 7, claim N14
(e) Chinese psychological operations in Taiwan, the South China Sea or elections	Cases are mostly single official claims or vendor reports with little traceability	Open	Section 6; Annex B incidents I04, I08, I19, I32
(f) Russian information operations outside Ukraine	France and Germany attribute publicly; the Romanian annulment remains disputed	Holds	Section 6; Annex B incidents I13, I17, I34
(g) NATO and US false flags in Europe (reconstructed family)	Stay-behind networks existed; courts found far-right perpetrators; NATO or CIA direction is not established	Adjusted (thesis bundles three propositions)	Section 7; Annex B rows B1 to B5 and N8

Source: Authors' compilation from Sections 6 and 7 and Annex B, using [116, 153, 175, 180, 234, 237, 238, 253, 272, 296]. Grades are claim-level and not event-level; the SARS-CoV-2 origin question in family (c) concerns accident versus natural spillover and not deliberate staging.

Two patterns run through the log. The sources that verified cleanly were declared doctrine and indexed journals, and the sources most exposed to error were short, regional or secondary items. Among the narrative families, claims were settled only where an institution could reach the scene, which is the access rule that Section 7 states.

Annex D. Defensive operating protocols

Section 8 ended on a rule and a procedure. The rule says that a claim of staging and a claim of authorship each need a stated threshold of evidence (Table 23). The procedure tells an analyst what to ask, and it leaves three questions open: who keeps the evidence and how, in what order the work is done, and what an organization may say before the answer is known. This annex answers them. D.1 sets out forensic protocols for keeping and testing digital material. D.2 turns the procedure into a guide for the first 24 hours, 72 hours and 30 days. D.3 organizes the response in three parallel flows, built so that a hypothesis of staging never becomes a public accusation ahead of the evidence. D.4 gives the record templates, and D.5 states the limits.

Two kinds of reader can use it. An affected organization owns the systems, accounts or sites involved and can contain, log and preserve. An outside observer, such as a research body or a newsroom, sees only public traces and cannot touch the underlying systems. Steps that need system access are marked for the organization, and the rest apply to both. An incident here means an event, or a claim about an event, that may involve a deniable operation: an intrusion, an act of sabotage with digital traces, or a circulating claim of responsibility or of staging.

The sources are standards and legal texts: the Berkeley Protocol [285], RFC 3227 [404], NIST SP 800-86 [405] and SP 800-61 Rev. 3 [403], ICD 203 [281] and the clocks of EU and Spanish law. These are Tier 1 material. Their combination into roles, gates and thresholds is the authors' proposal (Tier 3), and it has not been tried on a live incident. No source read supplies error rates for the indicator classes of Table 22, so nothing here claims that following the routines makes a judgment correct. The claim is narrower. The routines make a judgment auditable, so that a second person can find the error.

The annex is defensive. It describes how to keep, test and report evidence. It contains no guidance on staging an incident, disguising authorship or evading detection, and where a technique cuts both ways, as infrastructure analysis does, it states only what a defender needs in order to avoid a wrong conclusion. Three limits hold throughout. Collection is passive. No attribution rests on a single overlap. A claim that an incident was staged needs the same evidence as a claim that a named actor carried it out.

Table 59. Annex D at a glance

Part	Question it answers	Main sources	What it produces
D.1 Forensic protocols	How is digital material kept, authenticated and tested so that others can audit it?	Berkeley Protocol; RFC 3227; NIST SP 800-86	Custody log, hash list, test grades, discarded-hypotheses log
D.2 Verification guide	What is done, kept and said in the first 24 hours, 72 hours and 30 days?	NIS2; GDPR; Regulation 2023/1543; ICD 203	Phase steps, legal clocks, statement ceilings
D.3 Response flows	Who decides, in which order, and what stops an accusation running ahead of evidence?	NIST SP 800-61 Rev. 3; Table 23	Roles, decision gates G1 to G4, symmetry rule
D.4 Templates	Which records make the work auditable?	Authors' proposal	Six records (four new tables, Tables 76 to 79)
D.5 Limits	What does the annex not do, and where does it meet the manual?	Sections 8 and 11	Limits and reading route

Source: Authors' proposal, Tier 3, built on [281, 285, 403, 404, 405].

Key finding

The annex adds no new analytic test. Each gate in D.3 reuses the three tests of Section 1 (T, C and I), the three hypotheses of Table 23 and the five-part statement of Section 8. What it adds is who does what, in which order and with which record (Tier 3).

D.1 Forensic protocols

Forensic work, as this manual uses the term, means keeping digital material in a state in which a stranger can check what it is, where it came from and what has been done to it. The standards agree on the core. The Berkeley Protocol asks that every step, "from identification of relevant material through collection, analysis and reporting", be documented, because anyone who handles online information should expect the method to be questioned (para. 25) [285]. RFC 3227, the oldest text used here, tells the person handling a compromised system to do "collection first and analysis later" (section 2) [404]. NIST SP 800-86 splits acquisition into a plan, the acquisition itself and a check on integrity (section 3.1.2) [405]. ISO/IEC 27037:2012 covers the identification, collection, acquisition and preservation of digital evidence; only its informative sections are public, so this annex uses its scope and definitions and leaves its clause text aside [406]. None of these texts was written for the attribution of hybrid operations. The Berkeley Protocol was written for international crimes (Section 8), and the other three for administrators and incident handlers. Using them here extends their method (Tier 3).

One rule frames all six protocols: collection is passive. The defender collects what it already holds, such as its own logs, mailboxes and server records, and what third parties publish. It does not scan, probe, log in to or fetch content from systems it suspects, and it does not interact with the accounts behind a claim. The reason is practical and legal at once. Contact can alert an operator, alter the evidence and expose the defender's own network, and it raises legal questions that differ by jurisdiction and that only counsel can answer. RFC 3227 asks for the backing of the site's established procedures and the privacy rules of the jurisdiction, and suggests that law enforcement confirm the procedures are adequate (sections 1 and 2.3) [404]. The Berkeley Protocol asks investigators to understand the laws that apply, data protection law included (para. 28) [285]. The data protection duty is concrete: personal data collected along the way must be adequate, relevant and limited to what is necessary, and kept no longer than needed (GDPR Article 5(1)(c) and (e)) [407]. If direct viewing of a suspect site becomes unavoidable, it is a decision for the Incident Lead with legal approval, taken from an isolated environment and logged like any other handling (Tier 3).

D.1.1 Digital chain of custody

A chain of custody is "the chronological documentation of the sequence of custodians" of a piece of evidence, together with the control, date and time, transfer, analysis and disposition of each item (Berkeley para. 167) [285]. RFC 3227 turns that into four questions the record must answer. Where, when and by whom was the item found and collected? Where, when and by whom was it handled or examined? Who had custody, during what period, and how was it stored? When custody changed, when and how did the transfer occur (section 4.1) [404]? Table 60 gives the fields of a log that answers them.

Table 60. Fields of a digital custody log

Field	What the entry must hold	Basis
Item identifier	A unique number, never reused, carried by every copy of the item	Berkeley para. 162
Origin and method	Where the item was found, how it was captured, with which tool or query	RFC 3227 sec. 4.1; Berkeley para. 155
Collector and time	Name, role, time in UTC and the measured offset of the system clock	RFC 3227 secs. 2, 3.2; Berkeley para. 155
Integrity value	Hash value with the algorithm named, computed at the moment of collection	Berkeley para. 155; SP 800-86 sec. 3.1.2
Storage	Location, medium and the people who hold access or keys	RFC 3227 sec. 4.2; Berkeley para. 170

Field	What the entry must hold	Basis
Access and handling	Every access: who, when, with what tool and for what purpose	RFC 3227 sec. 4.1; Berkeley para. 167
Transfers	Sender, receiver, time and method, with the hash checked at both ends	RFC 3227 sec. 4.1; SP 800-86 sec. 3.1.2
Derived copies	Each working copy or conversion, with tool, version and settings	Berkeley paras. 169, 197

Source: Authors' proposal, Tier 3, built on [285, 404, 405]. The fields combine requirements that the three texts state separately.

Not every item needs the full log. NIST SP 800-61 Rev. 3 notes that formal chain-of-custody handling "might not be performed for every incident" and still treats collected incident data as evidence (RS.AN-07) [403]. SP 800-86 sets the default: if it is unclear whether evidence needs to be preserved, it generally should be (section 3.1.2) [405]. The proposal here is a two-track rule decided at intake. Any item that may support an attribution-related statement, a regulatory filing or a legal claim goes on the formal track with every field of Table 60, and other items carry at least an identifier, a source, a time and a hash. The Evidence Custodian (D.3) makes the call, and doubt resolves toward the formal track (Tier 3).

Time deserves a line of its own. RFC 3227 asks the handler to note the difference between the system clock and UTC and to say, for each timestamp, which one is used (section 2), and the Berkeley Protocol asks that the collecting machine's clock be synchronized with a Network Time Protocol server so that time-related metadata are accurate (para. 155) [285, 404]. A log kept in local time, or on clocks of unknown drift, cannot be matched reliably against logs kept by others (Tier 3).

The person who captures an item should not be the only keeper of its record. SP 800-86 suggests that, where possible, one person at the scene be designated the evidence custodian, with sole responsibility for photographing, documenting and labeling each item and recording each action, who did it, where and when [405]. The proposal adds a second keeper: a copy of the hash list is held by another person, so that no single person can change both an item and the record that vouches for it (Tier 3).

D.1.2 Preservation of evidence

Online material is perishable. Platforms remove content under their terms of use and users edit or delete their own posts, so the Berkeley Protocol warns that online information "can be easily decontextualized, lost, erased or corrupted" (para. 157). It asks investigators to assess the likelihood of removal, and when removal is likely to

collect the most reliable known version at once and keep looking for earlier or better versions afterwards (para. 150) [285].

How much to capture depends on the use. A screenshot or a PDF may do for internal work product, but content with probative value calls for a sounder method (para. 153). The Protocol treats three elements as the minimum standard for evidence in court, the URL, the HTML source code and a full-page capture, while noting that such standards vary by context. Best practice adds embedded media, embedded metadata, surrounding context, collection data and a hash value (para. 155) [285]. For a compromised system, RFC 3227 adds an order. Evidence is collected from the most volatile source to the least, from registers and memory through temporary file systems and disk to remote logs and archival media (section 2.1). The handler makes a bit-level copy of the media and does forensic work on a copy of that copy, because analysis alters access times (section 2) [404]. SP 800-86 asks the handler to rank sources by likely value, volatility and effort, and to verify integrity by computing a digest of the original and of the copy and comparing them (section 3.1.2) [405].

Table 61 sets out eight controls, each tied to the failure it prevents.

Table 61. Preservation controls and the failure each prevents

Control	Rule	Failure prevented	Basis
Evidentiary copy	Keep a clean original in every format in which the item was collected	Later edits mistaken for the original	Berkeley para. 168
Working copy	Analyze copies only, stored apart from the evidentiary copy	Changes to access times and content	RFC 3227 sec. 2; Berkeley para. 169
Hash at collection	Compute SHA-256 at once and name the algorithm in the log	Silent change; disputed integrity	Berkeley para. 155; SP 800-86 sec. 3.1.2
Immutable storage	Write-once or retention-locked medium, with access limited and recorded	Deletion, overwriting, insider alteration	RFC 3227 sec. 4.2; SP 800-61r3 RS.AN-06
Second hash list	A copy of the hash list held by a second person	One person changing file and record	Authors' proposal
Fixity checks	Re-hash a random sample of stored items on a fixed schedule	Silent degradation of stored files	Berkeley para. 172
Backup	Three copies, two kinds of storage, one copy held elsewhere	Loss through failure or disaster	Berkeley para. 171
Storage audit log	Export the store's access events to a second system	Undetected or undocumented access	RFC 3227 sec. 4.2; SP 800-61r3 RS.AN-06

Source: Authors' proposal, Tier 3, built on [285, 403, 404, 405]. Immutable storage and the second hash list are the authors' additions to what the texts require.

The choice of hash is not left to habit. The Berkeley Protocol tells investigators to pick the hash by the currently accepted standard and points to NIST (para. 155 and note 142) [285]. NIST announced on December 15, 2022 that SHA-1 should be phased out by December 31, 2030 in favor of the SHA-2 and SHA-3 groups [408], and RFC 3227 names "sha1sum" among its example tools, a reminder of how old the text is [404]. This annex therefore proposes SHA-256, a member of the SHA-2 group, with the algorithm named beside every value so that a later reader knows what to recompute. A second algorithm may be added when a receiving party requires one (Tier 3).

Retention is a decision made in advance. SP 800-61 Rev. 3 asks organizations to retain incident evidence under their own preservation procedures and retention policies, weighing the chance of prosecution against the cost of keeping the data and the means of reading it (RS.AN-07) [403]. Personal data in the store follow data protection limits on purpose and storage [407], so counsel and the data protection officer set the period before the first capture and not after the last.

Caution

A hash fixes a file's state from the moment it is computed. The Berkeley Protocol describes it as confirmation that content has not been modified since collection (para. 155) [285]. It cannot show that the file was genuine before then. That burden falls on the records of how, where and by whom the item was collected (Tier 3).

D.1.3 Metadata authentication

Section 8 warned that a timing mismatch discriminates only after the metadata itself has been authenticated (Table 22). Authenticating metadata means deciding which of five states a file is in, because each state supports a different conclusion (Table 62). The Berkeley Protocol supplies the main distinctions. Metadata can be created by the user, by other users, by a service provider or by any device that handles the file, and some are stripped before or during upload or by social media applications. Platforms often transcode uploaded media for playback, and then "the metadata will be a reflection of the new file, not the original". Where metadata have been stripped, the Protocol says, researchers should seek other ways to verify the item (para. 184) [285]. When the integrity of a file matters, it suggests digital forensic examination of the file itself (para. 183) [285].

Table 62. Five states of a file's metadata and what each can support

State	What it means	Test	What it can support
Original	Fields come from the capturing device or software; no sign of re-encoding	Compare copies from several platforms; check fields against content and anchors	A timing or provenance finding, if independently anchored
Altered	Values were edited after creation, by a tool or by hand	Look for fields that contradict each other or the content; compare earlier copies	A finding that the file was changed; it says nothing about who changed it
Stripped by a platform	Fields were removed on upload or by the application	Compare with a control file that passed through the same platform	Nothing either way; absence is no evidence of tampering
Reconfigured on upload	File was re-encoded; fields describe the new file, not the original	Locate an earlier or less processed copy; document the transformation	Dates of the derivative only, not of the event
Unresolved	The state cannot be determined from the copies available	Record why; keep searching for earlier copies	A lead for further collection, never a finding

Source: Authors' proposal, Tier 3, built on [285, 409]. A control file is a file of known origin and known metadata that is passed through the same platform from the analyst's own account, so that the analyst can see which fields that platform removes or rewrites.

Provenance has its own limit. The Protocol prefers "first copy found online" to "first copy online", because the original may have been removed and because closed channels such as email and private messaging groups may have carried the item before it appeared in public (para. 178) [285]. The same caution applies to dates. An upload date is a fact recorded by the platform, and a date of capture is a claim carried inside the file. The two should never share a column in the log (Tier 3).

Provenance standards add a signed layer where creators use them. Content Credentials are cryptographically signed metadata that describe a file's history. The joint information sheet of the US National Security Agency and its Australian, Canadian and British counterparts stresses that they supply context for judging authenticity and do not decide whether content is true. Any change to the signed portion invalidates the signature, and lack of provenance information "should not automatically make media less trustworthy", since provenance data are generally opt-in and may be withheld for legitimate reasons [409]. Absence therefore says nothing about manipulation. A valid credential is a claim by its signer about the file's history, and it is graded like any other source claim (Tier 3).

Caution

Grade the state of the file before the meaning of any field. A date read from a re-encoded file describes the copy, and treating it as the date of the event turns a platform artifact into evidence (Tier 3).

D.1.4 Geolocation and chronolocation

Geolocation is the identification or estimation of the place of an object, of an activity or of the point from which an item was generated. Chronolocation is the corroboration of the dates and times of the events a piece of imagery shows; the Protocol's own example is estimating the time of day from the length of shadows "along with other indicators" (paras. 190, 191) [285]. Both claims rest on anchors, features of the scene that can be matched to independent records.

The Protocol points to what to look for: buildings, flora and fauna, symbols and insignia (para. 188), "objectively verifiable information" such as the weather on a given day or the location of a building (para. 189), and spatial analysis that checks landscape features against satellite or other images, geodata and maps (para. 203) [285]. Table 63 sorts these into anchor types, each with an independent check and a typical weakness. The rule of this annex is simple (Tier 3). One anchor makes a lead. Two independent anchors make a supported location or time. An anchor that rests on the same source as the item does not count as independent.

Table 63. Anchors for geolocation and chronolocation

Anchor	Example	Independent check	Typical weakness
Fixed structures	Buildings, bridges, towers, skyline	Match against maps and satellite or aerial images	Look-alike structures; images of different dates
Terrain and vegetation	Ridge lines, coastline, tree species, crop stage	Compare with geodata and seasonal records	Season of the image differs from the reference
Signs, language, insignia	Road signs, shop names, uniforms, plates	Check records of where and when such signs were used	Signs can be reused, moved or imitated
Sun and shadow	Shadow length and direction, height of the sun	Compare with the sun's position for the claimed place and date	Needs known object height and camera angle
Weather	Rain, snow, fog, wet ground	Archived observations from a meteorological service	Local conditions differ from station records

Anchor	Example	Independent check	Typical weakness
Visible infrastructure	Road layout, rail lines, poles, street furniture	Compare with planning records and street-level images	Infrastructure changes over time
Second source	Other footage or photos of the same event	Independent captures from other angles and times	Same original re-shared; circular reporting

Source: Authors' proposal, Tier 3, built on [285]. The anchor types follow paras. 188, 189 and 203 of the Protocol; the checks and weaknesses are the authors' reading.

Chronolocation needs the same discipline. A shadow gives a time window only if the height of the object, the camera angle and the place are known, and a weather record gives a window only if a station is near enough to represent local conditions. A claimed time must fit all the anchors at once, and the file's own metadata count as one more anchor and never as the arbiter (Table 62). The Protocol names two traps. Content analysis suffers from circular reporting, decontextualization and misinterpretation (para. 187), and a clip with gaps may still be probative although the gaps affect the weight it can bear (para. 192) [285]. Identifying individuals from facial or other bodily features lies outside this annex, because the Protocol says it requires a forensic expert (para. 188 and note 159) [285].

Caution

The two-anchor rule is a convention and no validated threshold. No study located measures error rates for geolocation or chronolocation, which is why question 1 of Table 30 stays open (Tier 3).

D.1.5 Infrastructure analysis

Infrastructure analysis asks whether two pieces of online infrastructure belong to the same operator, and how strongly the public record says so. It is the technical core of the T test of Section 1, and it is where superficial coincidence does the most damage. The public record has layers, each with its own rules (Table 64). In the EU, domain registration data are governed by NIS2 Article 28. Registries and registrars must keep accurate and complete data in a dedicated database. Only data that are not personal data must be made public (Article 28(4)), and access to the rest runs through lawful and duly substantiated requests, which must be answered within 72 hours (Article 28(5)) [303]. The RDAP standard defines how registration data are returned. A record carries dated events, such as registration, last change and expiration, and the standard says that events can be future dated, so a date in such a record is not always a past event (RFC 9083, sections 4.5 and 10.2.3) [410]. TLS certificates are recorded in

public logs under the Certificate Transparency design, which RFC 9162 specifies, as an Experimental protocol, so that anyone can audit the activity of certification authorities [411]. Historical resolution records and page archives are collections held by third parties, and they are graded like any secondary source.

Table 64. Infrastructure layers and what an overlap can show

Layer	Public record	What an overlap can show	Benign explanations
Domain registration	Registry or registrar data; RDAP events	Common registration time, registrar or contact data	Bulk registration; privacy services; resold or re-registered domains
Name servers and DNS	Current and historical resolution records	Shared name servers or address history	Shared hosting and managed DNS serve many unrelated customers
IP addresses and networks	Regional registry records; routing data	Same network block or provider	Cloud and content networks serve unrelated customers
TLS certificates	Public certificate logs	Same issuance pattern or certificate names	Many unrelated sites receive certificates from the same authorities
Site content and code	Archived pages, source code, embedded identifiers	Reused templates, code or tracking identifiers	Copied templates; open-source kits; deliberate imitation
Timing	Dated events across all layers	Activity that clusters in time	Coincidence; a shared news cycle; automated schedules

Source: Authors' proposal, Tier 3, built on [303, 410, 411]. The column of benign explanations is the authors' reading and is not exhaustive.

A shared registrar, name server, hosting provider or certificate authority describes a commodity service used by many customers. Registration in the same hour describes a bulk-registration habit as readily as a single operator. The question for each overlap is how many unrelated parties could have produced it. An overlap also cuts both ways. What a defender can see, another party can see and copy, so an overlap supports H1 and H2 alike unless the log can name evidence that an imitator could not easily reproduce, such as a seizure, a court record or a multi-state technical record (Section 6). The grading rule follows (Tier 3). No single overlap grades above P on T. A grade of Y on T needs at least two independent layers that are unlikely to coincide through shared services, and temporal correlation alone never counts. For each overlap the log records the benign explanations considered and why they were set aside.

Two process rules keep the work defensive and auditable. First, record the time and source of every lookup, since registry data and resolution answers change and a lookup without a time cannot be compared with another. Second, keep infrastructure findings apart from attribution to a person or a State. That two domains share an operator is a technical finding at the level of T. It becomes an attribution-related statement only by passing Gate 4 (D.3) (Tier 3).

Key finding

Infrastructure overlaps are cheap to find and cheap to misread. The annex grades them conservatively on purpose: one overlap is a lead, two independent layers can support P or Y on T, and no overlap by itself separates a named actor from an imitator (Tier 3).

D.1.6 Quality control

Quality control makes the other protocols trustworthy, and the evidence for it is useful and limited. The Berkeley Protocol says that accuracy is "often improved through the use and testing of multiple working hypotheses and/or peer review" (para. 30), and it asks investigators to develop and apply multiple working hypotheses without favoring any theory (para. 27) [285]. Marcoci and colleagues (2019) tested the ICD 203 criteria themselves. Fifteen raters scored 64 analytic reports written for hypothetical problems, and despite training and calibration the reliability of individual ratings was poor, while aggregated ratings were good to excellent. The authors advised groups of at least three raters for reliable quality control [412]. The study measured agreement on the quality of reasoning and not the accuracy of attribution, and its problems were simpler than real ones. It supports the design of a review, and it says nothing about error rates.

Table 65 turns this into six steps. The design puts three people in front of any attribution-related statement: Analyst A, Analyst B, who replicates the tests from the evidentiary copy without seeing A's conclusions, and a Challenge Reviewer who sits outside the incident chain and argues against the leading hypothesis. The Berkeley Protocol's note on bias lists expectation, confirmation, anchoring and contextual bias among the effects that damage investigative quality (note 24) [285]. Blind replication limits anchoring on a colleague's conclusion and exposure to context that invites it (Tier 3). ICD 203 gives the independence principle a precedent, since its Analytic Ombuds reports directly to a senior official, holds no other position in the office and addresses concerns about objectivity (section E.2) [281].

Table 65. Six quality-control steps

Step	Who	Record	When
Blind replication	Analyst B, from the evidentiary copy	Replication note listing every difference from Analyst A	Before G3 and G4
Challenge review	Reviewer outside the incident chain	Memo arguing the best alternative to the leading hypothesis	Before G4
Method review	Senior or external reviewer	Sign-off on procedure and tool settings, not on the conclusion	Before a technical report is published
Discarded-hypotheses log	The analyst who drops a hypothesis	Hypothesis, reason, evidence, date and who agreed	At every rejection
Tool and version log	Whoever runs the tool	Tool, version, settings and input hash for each step	At every processing step
Re-grade on new evidence	Analyst A, replicated by Analyst B; Incident Lead approves the status change	Dated change of grade with the reason	When a new source or official finding appears

Source: Authors' proposal, Tier 3, built on [281, 285, 404, 412]. The three-person rule follows the minimum group size in Marcoci et al., a finding about rating reports on hypothetical problems.

The tool log has an old rationale: RFC 3227 asks the handler to be prepared to testify to the authenticity and reliability of the tools used (section 5), and the Berkeley Protocol asks investigators to record any tools or software used and to understand any automated process well enough to explain it (para. 25) [285, 404].

Key finding

An unreproducible finding is a lead. If Analyst B cannot reproduce a test from the evidentiary copy and the custody log, the test is graded no higher than P, however confident Analyst A is (Tier 3).

D.2 Verification operations guide

Section 8 offered a seven-step procedure and a table of thresholds. This part says what happens when, in the three windows in which an incident is usually decided: the first 24 hours, up to hour 72 and up to day 30. The windows are not arbitrary. They follow the clocks of EU law (Table 66), and they match an investigation in which preservation comes first, assessment second and attribution-related judgment last. The third window is set at 30 days to leave room for the NIS2 final report, which is due not later than one month after the 72-hour notification [303].

Windows are clocks, but statements are governed by gates. Time passing never entitles anyone to say more. A statement is allowed when its gate in D.3 is open, and a window that ends without evidence changes nothing (Tier 3).

Table 66. Legal and procedural clocks that shape the windows

Clock	Who must act	Deadline	Source
NIS2 early warning	Essential or important entity, after a significant incident	Within 24 hours of becoming aware	NIS2 Art. 23(4) (a)
NIS2 incident notification	The same entity	Within 72 hours of becoming aware	NIS2 Art. 23(4) (b)
NIS2 final report	The same entity	Not later than one month after the notification	NIS2 Art. 23(4) (d)
CSIRT first response	CSIRT or competent authority	Where possible within 24 hours of the early warning	NIS2 Art. 23(5)
Spanish national guide	Obligated subjects under Royal Decree-Law 12/2018	Initial at once; follow-ups by impact level, up to 40 days	INCIBE guide (2020), its Table 7
GDPR breach notice	Controller of a personal data breach, unless risk is unlikely	Within 72 hours of becoming aware, where feasible	GDPR Art. 33(1)
Registry access reply	Domain registries and registrars	Within 72 hours of a lawful, substantiated request	NIS2 Art. 28(5)
EU preservation order	Service provider addressed by a European Preservation Order	60 days, with one extension of 30 days	Reg. 2023/1543 Art. 11(1)
Spanish preservation order	Any person or company, on request of the prosecutor or judicial police	90 days, extendable once, to 180 days at most	LECrim Art. 588 octies
Budapest preservation	Person holding the data, on order of a Party's authority	Up to 90 days, renewable if allowed; at least 60 days on request	Convention Arts. 16(2), 29(7)

Source: Authors' compilation from [303, 407, 413, 414, 415, 416]. NIS2 clocks bind entities through national law, and Spain had not transposed the Directive at the evidence cutoff (Table 26); the Spanish guide predates NIS2; the GDPR applies directly. Preservation orders are issued by judges, prosecutors or police and not by private organizations. The Budapest articles were read in the Council of Europe text, and the status of each State as a party was not checked.

Stat box

Four clocks set the pace: 24 hours for the NIS2 early warning, 72 hours for the NIS2 notification and for the GDPR notice, and one month after the notification for the NIS2 final report. Preservation orders in the texts reviewed hold data for 60 to 180 days, but only from the moment an authority issues one, so Table 69 sends third-party evidence to Legal within the first 24 hours.

In Spain the channels that exist today come from earlier law. The national guide on cyber-incident notification, approved by the National Cybersecurity Council on 21 February 2020, names four reference bodies: CCN-CERT for the public sector and for systems that handle classified information, INCIBE-CERT for citizens and the private sector, the CNPIC for critical infrastructure and critical operators, and ESP-DEF-CERT for the Armed Forces and other systems that affect national defense. Under its single-window system the affected party sends an email or ticket to its reference CSIRT (INCIBE-CERT or CCN-CERT), which passes the case on according to the incident: to ESP-DEF-CERT if national defense is affected, to the CNPIC if a critical infrastructure under Law 8/2011 is involved, and to the Spanish Data Protection Agency (AEPD) if the GDPR is engaged. Where appropriate, the CNPIC's Cyber Coordination Office then makes the information available to the State security forces and the Public Prosecutor so that a police and judicial investigation can begin [416]. For obliged subjects the guide sets an initial notification at once for incidents of critical, very high and high impact, an intermediate notification 24 or 48 hours (critical) or 72 hours (very high) after the initial one, and a final notification after 20 days (critical) or 40 days (very high) [416]. The guide predates NIS2 and rests on Royal Decree-Law 12/2018, and Spain's NIS2 transposition law was not yet in force at the evidence cutoff (Table 26) [313]. A reader should confirm the current channels, and the authority designated under the new law, before relying on any of this.

D.2.1 The first 24 hours: stabilize, preserve, hold

The aim of the first window is to stop the evidence from degrading while saying as little as the facts allow. Seven steps, in this order, apply to an affected organization.

1. Declare the incident and record when and why the organization became aware. The NIS2 and GDPR clocks run from awareness (Table 66), so the time and the reason go into the intake sheet (Table 76) before anything else.
2. Name the Incident Lead and the Evidence Custodian and open the custody log. NIST SP 800-61 Rev. 3 suggests designating an incident lead for each incident (RS.MA-01) [403], and SP 800-86 an evidence custodian [405].
3. Capture the items most at risk first, in the order of volatility of Table 67, hash each at once and store it as in D.1.2. Collection comes before analysis.
4. Contain without destroying. Prefer isolating a system to wiping or powering it down, export logs before they rotate, and choose the method of isolation with the investigators, because simply disconnecting or filtering a host from the network may trigger "deadman switches" that wipe evidence (RFC 3227, section 2.2) [404].

SP 800-86 asks that the containment decision follow existing policy and the team's assessment of risk, so that the strategy chosen mitigates the risk while maintaining the integrity of potential evidence whenever possible (section 3.1.3) [405].

5. Open the hypothesis log. Record the claim and its channel, note who could and who could not inspect the scene, the logs or the wreckage, and set out H1, H2 and H3 of Table 23 with, for each, the evidence that would count against it (Section 8, steps 1 to 4). Nothing from this log leaves the team.
6. Run the escalation test of Table 69 and send what the clocks require. The NIS2 early warning indicates, where applicable, whether the incident is suspected of being caused by unlawful or malicious acts and whether it could have cross-border impact. It does not ask who did it (Table 25) [303].
7. If outside attention is likely, issue a holding statement that passes Gate 1 (D.3): confirmed facts only, and the time of the next update.

An outside observer follows the same order with fewer levers. It archives the claim and its context before engaging with it, meaning URL, account, post, time and first copy found, hashes and logs the captures, and publishes nothing, a repost with commentary included, until Gate 1 is passed. Amplifying a claim before it is tested is itself a handling error (Table 68).

D.2.2 Up to 72 hours: test, replicate, assess

The second window turns captured material into graded findings and files the first regulatory notices.

1. Complete the acquisition. Take bit-level copies of affected media, collect remote logs and third-party records, and verify every hash after copying (SP 800-86, section 3.1.2) [405].
2. Authenticate and locate. Assign each file a metadata state (Table 62), look for a second anchor for every claim of place or time (Table 63), and map infrastructure by layer without contact (Table 64).
3. Grade T, C and I for H1 and H2 and apply the Table 23 criterion to H3, discounting evidence that fits all three hypotheses (Section 8, step 5), and log every hypothesis set aside, with the reason.
4. Replicate. Analyst B repeats the key tests blind from the evidentiary copy (Table 65).

5. File the notices that are due. For an entity within the scope of NIS2 and a significant incident, the incident notification is due within 72 hours of becoming aware, with an initial assessment of severity and impact and, where available, indicators of compromise [303]. If personal data are breached, the controller notifies the supervisory authority within 72 hours where feasible, gives reasons for any delay and may supply information in phases (GDPR Article 33(1) and (4)) [407].
6. Ask for third-party preservation. Orders to hold data at platforms, hosts and carriers come from judges, prosecutors or the police and not from private organizations (Table 66). The organization's task is to ask early, through Legal, and to hand over identifiers and dates.
7. Pass Gate 2 and, only if Analysts A and B concur, Gate 3, and release the first assessed statement in five parts (D.2.7).

An outside observer cannot file notices or ask for orders, but it can still replicate. A second analyst repeats the geolocation and metadata tests on the observer's own captures before any assessment is published.

D.2.3 Up to day 30: review, finalize, correct

The third window ends with a final record and leaves room for the final report.

1. Run the challenge review on the leading hypothesis and file the memo (Table 65).
2. Re-grade T, C and I after replication and review, and record every change of judgment with its reason (Section 8, step 7).
3. Prepare the final report: a detailed description of the incident with its severity and impact, the type of threat or root cause that is likely to have triggered it, the mitigation applied and any cross-border impact. If the incident is still ongoing at that point, a progress report is due and the final report follows within one month of the handling of the incident (NIS2 Article 23(4)(d) and (e)) [303].
4. Decide on Gate 4 only with Legal's sign-off. A statement that links the incident to an actor uses "consistent with" wording or cites an official attribution, and naming an actor needs the threshold of Table 23 (D.3).
5. Move the evidence to long-term retention under the retention policy and any legal hold, and run a fixity check on a sample (Table 61).
6. Hold the post-incident review and file the after-action report (D.3.5).
7. Set review triggers, such as a court filing, an authority's report or a vendor finding, and a date to re-check the statement register.

An outside observer treats day 30 as a publication checkpoint. It publishes a method note with the assessment, the scale and confidence used, and what would change them.

D.2.4 Evidence to preserve before it disappears

The order of capture follows volatility and not importance. RFC 3227 puts registers, caches, routing and process tables and memory first, temporary file systems and disk next, then remote logs, physical configuration and archival media (section 2.1) [404]. SP 800-86 names three factors for ranking sources, likely value, volatility and effort, and notes that non-volatile data can be somewhat dynamic too, as log files are when new events overwrite them (section 3.1.2) [405]. Public online material has a volatility of its own, set by platforms and by users (Berkeley paras. 150, 157) [285]. Table 67 lists the classes of evidence that matter in an incident that may involve a deniable operation, with what to capture and why each class is at risk.

Table 67. Evidence classes, what to capture and why each is at risk

Class	What to capture	Why it is at risk	Handling note
Pages and posts	URL, HTML source, full-page capture, embedded media, comments, collection data	Edited, deleted or removed by the poster or the platform	Capture first; hash; note the tool and time
Accounts and profiles	Handle, identifiers, profile text, posting history, links to other accounts	Renamed, deleted or locked	Capture identifiers as well as names
Files and media	Original file in native format; every copy found; the path it traveled	Re-encoding strips metadata; reposts overwrite	Keep the untouched original; work on copies
Domains and hosting	Registration data, DNS answers, certificates, archived pages, each with lookup time	Registrations lapse, records change, sites go dark	Record time and source of every lookup
Network and system logs	Firewall, proxy, DNS, authentication and application logs; volatile state where relevant	Rotation, overwriting, shutdown	Export before rotation; follow the volatility order
Communications	Messages and email with full headers; call and meeting records; internal chat	Deletion, retention limits, auto-delete	Preserve the full message source, not forwards or screenshots

Class	What to capture	Why it is at risk	Handling note
People and context	Who was present, what was seen and said, what the operators did	Memory fades and accounts converge	Note names, times and observations soon
Third-party records	Platform, hosting, carrier and registrar records	Held for limited periods	Ask Legal for early preservation requests

Source: Authors' proposal, Tier 3, built on [285, 404, 405]. The people-and-context row follows RFC 3227 section 3.2, which asks the handler to note who was present, what they did and what they observed.

Collection of communications and logs touches personal data, so the data protection officer is told before the first export, and the store follows the limits of purpose and storage that apply to personal data [407]. The sources pull in two directions. RFC 3227 advises erring toward collecting too much rather than too little (section 3.2) [404], while the Berkeley Protocol asks that content be collected only if it is justified for an articulable purpose, necessary to achieve it and proportional to it (para. 31) [285]. A shared inbox or a chat export may hold far more than the incident needs. A workable reading is to preserve whole sources, such as a mailbox export or a log set, under restricted access, and to analyze and share only the parts the question needs (Tier 3).

D.2.5 What premature handling contaminates

Most evidence is lost to helpful action and not to the adversary. RFC 3227 lists the usual ways: shutting down before collection is complete, trusting the programs on the compromised system, and running tools that change the access times of all files (section 2.2) [404]. The Berkeley Protocol adds that any alterations, transformations or conversions caused by the collection process should be documented (para. 155) and advises processing working copies only, never the original or evidentiary copy (para. 197) [285]. Table 68 lists the handling errors most likely in an incident of this kind and the safer alternative to each.

Table 68. Premature handling, what it damages and the safer alternative

Handling error	What it damages	Safer alternative
Rebooting or powering down	Memory, process state, temporary files	Capture volatile data first; then isolate by an agreed method
Wiping or rebuilding a system	Disk traces, malware samples, logs	Image the media; rebuild from a clean copy afterwards
Opening files in their native application	Access times, embedded content, metadata	Work on a hashed working copy in an isolated viewer
Forwarding evidence by email or chat	Headers, metadata and message source	Export the original with full headers into the evidence store

Handling error	What it damages	Safer alternative
Screenshots of screenshots, or reposting	Provenance and the first-copy trail	Capture the source page; keep the first copy found
Resetting passwords or deleting accounts before capture	Account state and login history	Capture logs and identifiers, then reset
Visiting or probing suspect sites from the organization's network	Evidence, the defender's exposure, the secrecy of the inquiry	Use archived copies and public records; ask Legal first
Naming an indicator or actor in public	The investigation's options and the credibility of later findings	Share through CSIRT or agreed channels; hold public detail

Source: Authors' proposal, Tier 3, built on [154, 285, 404, 405]. The last row reflects the costs of the public use of intelligence that Dylan and Maguire list (Section 9).

The last row deserves a word. Dylan and Maguire list the costs of making intelligence public: sources lost once exposed, narrowed options for escalation, and credibility damage when claims lack support (Section 9) [154]. A defender that publishes an indicator or a suspected actor pays the same costs on a smaller scale. NIS2 gives defenders a channel that avoids them. Article 29(1) lets entities exchange, on a voluntary basis, information that includes indicators of compromise and threat-actor-specific information, and Article 29(2) places that exchange within communities of essential and important entities, and where relevant their suppliers or service providers, under information-sharing arrangements [303]. NIST SP 800-61 Rev. 3 likewise recommends sharing information securely and in line with response plans and information-sharing agreements (RS.CO-03) [403].

D.2.6 When to escalate

Escalation fails in two ways. It comes too late, or it waits for an attribution that never arrives. Table 69 sets triggers that do not depend on knowing who acted. The NIS2 early warning asks only whether a malicious act is suspected (Table 25), and NIS2 Article 23(5) has the CSIRT or competent authority give guidance on reporting to law enforcement where a significant incident is suspected to be of criminal nature [303]. An organization in doubt can therefore put that question to its CSIRT. SP 800-61 Rev. 3 distinguishes escalation, which adds resources or time frames, from elevation, which brings in a higher level of management (RS.MA-04), and asks that notifications to law enforcement and regulators follow criteria in the response plan and management approval (RS.CO-02) [403].

Table 69. Escalation triggers

Trigger	Escalate to	When	Hand over
Any suspected malicious act	Cybersecurity lead and Incident Lead	At once	Custody log, hypothesis log, first captures
Personal data possibly involved	Legal and data protection officer	At once; the GDPR clock runs from awareness	Data categories, estimated numbers, containment steps
Significant incident in an entity within NIS2 scope	Legal, then national CSIRT or authority	Early warning within 24 hours	Suspected malicious act; possible cross-border impact
Signs of a crime such as extortion, fraud or sabotage	Law enforcement, through Legal	Without waiting for attribution	Preserved items with hashes and the custody log
Third-party evidence at risk of deletion	Legal, then police or prosecutor	Within the first 24 hours	Identifiers, platform names, dates, reason for risk
A hypothesis that a State acted	Competent national authority, through the CSIRT	Before any public wording	Indicators and gradings; no public attribution
Press, official or online inquiry on cause or actor	Communications Lead with Incident Lead	Before any reply	Current gate and approved wording
Analysts disagree on a grade	Challenge Reviewer, then Incident Lead	Before the next gate	Both grades with the evidence for each

Source: Authors' proposal, Tier 3, built on [303, 403, 407]. Triggers and timings other than the legal clocks of Table 66 are the authors' proposal.

In Spain the 2020 disinformation procedure has the National Security Council handle public attribution to a third State (Table 26) [308]. That text covers disinformation, and no cyber-specific attribution rule was found in the sources read, so the safe course for a private organization that suspects State involvement is to pass its indicators and gradings to its reference CSIRT and leave attribution to the authorities (Tier 3). When physical safety is at stake, the usual emergency routes come before anything in this annex.

D.2.7 What may be said in each phase

Table 70 gives the highest statement that is normally open in each window and what has to stay a hypothesis. Each row also carries a wording test that any reviewer can apply to a draft.

Table 70. Statement ceilings by window

Window	May be said	Stays a hypothesis	Wording test
Hours 0 to 24	An incident is being handled; what is affected; what users should do; next update	Cause, actor, motive, any claim of staging	Could a log or a person confirm each sentence today?
Hours 24 to 72	Confirmed technical facts; whether a malicious act is suspected; scope so far	Actor, State involvement, staging	Does each "suspected" say what would confirm it?
Days 3 to 30	Assessed cause and scope on a named scale, with separate confidence	Any named actor without official attribution or Table 23 evidence	Would the sentence survive Analyst B's replication?
After day 30	Final cause and corrections to earlier statements	Whatever stays below threshold, labeled with what would change it	Does it say what changed since the last statement, and why?

Source: Authors' proposal, Tier 3, built on [281, 303]. A ceiling applies only if the matching gate of Table 73 is open.

The ceilings are ceilings and no entitlements. NIS2 asks the entity to report to its CSIRT and, where appropriate, to notify the recipients of its services of significant incidents likely to affect them (Article 23(1)). It does not ask the entity to name a cause or an actor in public. Under Article 23(7) it is the CSIRT or competent authority that may, after consulting the entity, inform the public or require the entity to do so, where public awareness is needed to prevent or deal with a significant incident or disclosure is otherwise in the public interest [303]. Keeping the public statement within its ceiling is therefore compatible with full reporting to the authorities, and the two should be planned together (Tier 3).

D.2.8 Documenting uncertainty without feeding speculation

Uncertainty has to be written down in two places that serve different readers. The hypothesis log is the internal record and holds everything, including hypotheses the evidence does not yet support. The statement register holds only what passed a gate. Keeping the two apart is the main defense against speculation, since a hypothesis does not become a statement by being logged. Section 8 already supplies the grammar, a statement in five parts: what is known; what is assessed, with a named probability scale and a separate confidence level; what is not known; what would change the assessment; and when the next update will come.

ICD 203 supplies the discipline behind it. It asks analysts to explain the basis of their uncertainty as likelihood and confidence, to draw probability terms from one stated set, with a strong encouragement not to mix rows, and never to combine a confidence level and a degree of likelihood in the same sentence (standard 2). It asks them to

separate information from assumptions and judgments (standard 3), to assess plausible alternatives (standard 4), to name indicators that would alter a judgment (standards 2 to 4) and to say how a judgment has changed or stayed the same (standard 7) [281]. The UK yardstick offers another vocabulary, and whichever is used must be named [101]. Six working rules follow (Tier 3).

- Say what is known in sentences that a log or a person can confirm, and put assumptions in separate sentences labeled as such.
- State the probability on one named scale and the confidence separately.
- Replace "we are still investigating" with the event that would change the assessment, such as a forensic result, a court filing or an authority's report.
- Keep "unknown" neutral. List what is unknown without hinting at who or why.
- Give the time of the next update and keep it, even when nothing has changed.
- Correct visibly. When a judgment changes, say what changed and why, in the same register as the original statement.

Spain supplies a test case of speculation running ahead of evidence. On 29 April 2025, the day after the Iberian blackout, an Audiencia Nacional judge opened proceedings on possible cyberterrorism and the prime minister said no hypothesis was ruled out [277]. The government committee reported on 17 June 2025 that no evidence existed of a cyber incident as the cause, and on 12 January 2026 the court provisionally dismissed the case for lack of any minimal indication of terrorist sabotage (Section 7) [278, 279]. *Inference*. Keeping a hypothesis open was defensible. A five-part statement would have kept it open without offering a cause: it would have listed what was known, said that no cause was yet assessed, named what was unknown and said that the technical inquiry was the event that would change the assessment (Tier 3).

Caution

A denial, a retraction and a hypothesis each repeat the claim they address. Wording that names a suspected actor in order to say it is not yet attributed still carries the name. Prefer "the actor is not yet identified" to "we cannot yet say whether X was responsible" (Tier 3).

D.3 Incident response in three parallel flows

A response to a suspected deniable operation is three jobs and not one. Someone has to stop the damage, someone has to find out what happened, and someone has to tell people about it. The three move at different speeds. Containment answers to operations, investigation answers to evidence, and communication answers to the

press, the regulator and the public. When they pull apart the fastest usually wins, and the fastest is almost always communication. This part shows how the flows stay in step: what each one does, who runs it, what passes between them, and the four gates that let a statement out only when the evidence behind it exists.

D.3.1 Three flows, one clock

Table 71 sets out the flows. The first three columns are the ordinary content of an incident plan. The last column is the addition that matters here, a plain statement of the mistake each flow makes when it runs alone.

Table 71. The three flows of an incident response

Flow	Objective	Example actions	Must not do
Containment	Limit damage and avoid destroying evidence	Isolate systems by an agreed method; export logs; block malicious domains	Wipe, reboot or probe suspect systems before capture
Investigation	Establish facts, scope and possible causes	Forensic analysis; access review; test indicators; grade T, C and I	Let a favored hypothesis steer what is collected
Communication	Protect audiences and avoid over-attribution	Holding statements; graded messages; notices to regulators	Accuse, or hint at an actor, before a gate is open

Source: Authors' proposal, Tier 3. The flows correspond to the mitigation (RS.MI), analysis (RS.AN) and reporting and communication (RS.CO) functions of [403]; the last column is the authors' addition.

NIST SP 800-61 Rev. 3 draws the same lines. It separates incident mitigation, incident analysis and incident response reporting and communication, and it places incident management (RS.MA) above them as the layer that coordinates [403]. The standard is frank about the friction between the first two. It asks teams to choose a strategy by balancing the need to recover quickly against the need to observe the attacker or investigate more thoroughly, and it adds that every such decision has trade-offs (RS.MA-03) [403]. Where a team considers delaying containment in order to watch an intruder, it should first discuss the feasibility with the legal department, and it should know that the delay can be dangerous because the attacker could escalate access or compromise other systems (RS.MI, R1) [403]. The annex takes the same line. A decision to hold back containment for the sake of evidence is made by the Incident Lead together with Legal, written in the intake sheet with its reason, and reviewed at fixed intervals (Tier 3).

Investigation has a trap of its own. The standard asks the team to determine the sequence of events and the assets involved, to find the underlying root causes, and to

"attempt to determine" which vulnerabilities, threats and threat actors were involved (RS.AN-03, R1 to R3) [403]. The wording is modest: an attempt, and one item among three. NIS2 is narrower still, because the final report asks for the type of threat or root cause that is likely to have triggered the incident, and not for an actor (Article 23(4)(d)) [303]. The annex therefore treats cause and actor as two questions with two budgets. The cause is what the investigation owes the regulator and the organization. An actor is reported when the evidence reaches the threshold of Table 23, and is otherwise left open (Tier 3).

Communication is not a single channel either. SP 800-61 Rev. 3 sorts reporting and communication into four kinds: coordination among the parties with response roles, notification of affected parties and regulators, public communication such as answering media inquiries, and the voluntary sharing of threat information (RS.CO, note 1) [403]. Only the third faces the general public. The other three carry more detail to fewer people, and each keeps its own record. The External Liaison logs coordination and sharing, Legal logs notifications, and the Communications Lead keeps the statement register for public statements (D.4). The gates of D.3.3 bind every channel in one respect: an assessment of cause goes out only after Gate 3, and a link to an actor only after Gate 4. Raw indicators can travel earlier in trusted channels, with their source, time and grade, because the readers there can test them and are bound by an information-sharing agreement (Tier 3).

Three handovers keep the flows honest.

- Containment hands the investigation a log of everything it changed, with times and reasons, so that analysts can tell original traces from consequences of the response.
- Investigation hands communication graded findings through the gates and never a raw hypothesis.
- Communication hands the investigation every outside claim, question and third-party report, logged as a claim with its channel (Section 8, step 1) and not as evidence. The standard asks that third-party reports be monitored and taken seriously (RS.MA-02, R2) [403].

D.3.2 Roles and independence

A response team that is small, tired and under pressure fails in predictable ways, and most of them put two incompatible jobs on one person. The analyst who found the lead also grades it. The person who drafts the statement also decides how far the

evidence goes. Table 72 separates eight roles so that nobody marks their own work. SP 800-61 Rev. 3 suggests designating an incident lead for each incident (RS.MA-01) [403], and SP 800-86 an evidence custodian [405]; the other independence rules are the authors' proposal.

Table 72. Eight roles and the independence rule for each

Role	Responsibility	Independence rule
Incident Lead	Runs the response across all three flows; opens Gates 1 to 3	Does not grade evidence or replicate tests
Evidence Custodian	Keeps the custody log, the evidence store and the hash lists	Reports integrity problems straight to the Incident Lead and Legal
Analyst A	Leads the analysis, keeps the hypothesis log and grades T, C and I	Does not draft public statements
Analyst B	Repeats the key tests blind from the evidentiary copy	Does not see A's grades until the replication note is filed
Challenge Reviewer	Argues against the leading hypothesis; reviews wording at Gate 4	Sits outside the analysis team
Legal and Data Protection	Advises on notices, preservation requests, privacy and wording; signs off Gate 4	Is not measured on speed of response
Communications Lead	Drafts statements within the open gate; keeps the statement register	Issues nothing that the register does not hold
External Liaison	Handles CSIRT, authorities, platforms and peers; keeps the contact log	Shares indicators only through agreed channels

Source: Authors' proposal, Tier 3, built on [403, 405]. In a small organization one person may hold several roles, subject to the three exceptions in the text.

In a small organization one person may hold several roles. Three pairs should never merge: Analyst A and Analyst B, the analysis team and the Challenge Reviewer, and the Communications Lead and whoever opens the gates (Tier 3). An organization without a second analyst can borrow one. NIS2 Article 29 allows voluntary exchange of cybersecurity information within communities of entities, under information-sharing arrangements [303], and the Berkeley Protocol notes that where partnerships are formed it can be helpful to establish a written agreement for information-sharing (para. 118), and that memorandums of understanding or contracts should ensure partners comply with the sharing policy (para. 131) [285]. A borrowed reviewer receives the evidentiary copy, the custody log and the hypothesis log, and does not receive the leading conclusion until the replication note is filed.

An outside observer such as a newsroom, a university group or a civil-society monitor needs fewer people and the same separations. Three will do: one who holds the evidence and opens the gates, one who analyzes, and one who replicates and challenges, with a partner desk standing in for any role the team cannot fill.

D.3.3 Decision gates

A gate is a checkpoint in the stream of statements. It has an opening condition, a record that proves the condition was met, and a ceiling on what may be said once it is open. Gates are taken in order, but they open on evidence and never on elapsed time. Day 3 does not open Gate 3, and a quiet week opens nothing. They can also close. If a replication fails or a new fact contradicts the basis of a statement, the gate that supported it closes and the next statement says so (Table 73).

Table 73. Four decision gates

Gate	Opens when	Record required	Highest permitted statement
G1 Preservation	First captures are hashed and logged, and the facts to be stated are confirmed	Intake sheet; custody entries	Confirmed facts, advice to users, time of next update
G2 Factual	Scope is confirmed by logs or by a second source	Evidence list with metadata states	Confirmed technical facts; that a malicious act is suspected, if facts support it
G3 Assessed	Analysts A and B concur; alternatives and disconfirming evidence are logged	Hypothesis log; replication note	Assessed cause on a named scale, with separate confidence
G4 Attribution-related	T, C and I graded and replicated; challenge memo filed; Legal signs off	Challenge memo; legal sign-off; register entry	A cited official attribution, or wording at the strength Table 23 supports

Source: Authors' proposal, Tier 3. The permitted statements apply the five-part form of Section 8 and the thresholds of Table 23. No gate has been tested against a live incident.

Four rules complete the table.

- Gates govern what an organization chooses to say, and never delay what the law requires. A notice due under NIS2 or the GDPR goes out on time and in the form the law asks for, stating what is known or suspected and labeling each as such (Table 66).
- The ceiling travels with the gate. A statement may say less than its gate permits and never more.

- At Gate 4, "consistent with" wording needs a stated similarity, such as a technique or an infrastructure overlap, and must say what the similarity does not show. Naming an actor as responsible needs an official attribution to cite or evidence at the Table 23 threshold for H1, which is two of T, C and I graded Y from independent sources.
- A private organization does not attribute to a State on its own account. It may cite an official attribution, and it passes its indicators and gradings to the authorities (D.2.6).

The three-person rule at Gate 4, Analyst A, Analyst B and the Challenge Reviewer, is the one set out in D.1.6, and it rests on the same indirect evidence.

Caution

Gate 4 is where an organization is most tempted to treat a legal sign-off as an evidential one. Legal clears wording and exposure. It does not raise the grade of the evidence, and a statement that Legal has approved is not for that reason better supported (Tier 3).

D.3.4 The symmetry rule for claims of staging

One rule governs every public statement about staging: a claim that an incident was staged or disguised to blame another party needs the same level of evidence as a claim that a named party carried it out. Section 8 sets the thresholds for analysts (Table 23). What an organization under pressure needs is the same rule turned into wording, and Table 74 does that.

The reasons run in both directions. Accusing a named party without evidence is the familiar error. The mirror error is declaring a false flag without evidence. The nearest experimental evidence is the liar's dividend: in five pre-registered experiments with 15,287 US adults, politicians who falsely called a real text-based scandal fake gained support, and the tactic was largely ineffective against video (Section 7) [267]. *Inference*. Calling a real incident staged is the same move, so declaring staging without discriminating evidence may hand the declarer the same dividend and cost it credibility if the claim fails; no study retrieved tests that analogy for false-flag claims (Tier 3). Ignoring deception has a documented cost as well. A 1972 CIA text records that analysts accepted the 1968 Czechoslovakia "exercises" at face value because earlier announced exercises had been real [288]. The rule is therefore no license for silence. Keeping H2 open costs little, and the hypothesis log is where it lives (Table 23 asks for one indicator, a third party with motive and disconfirming evidence named in advance). The rule governs the sentence that leaves the building.

Table 74. Evidence bars and wording for four claims

Claim	Minimum evidence to assert	Wording while below the bar	Wording not allowed
A named party acted (H1)	Two of T, C and I graded Y, from independent sources	Actor not yet identified; similarity of technique does not identify an actor	Naming a party; "all signs point to"; "who benefits?" framing
Staged or disguised to blame another (H2)	Tier 2 evidence of design: planning papers, a traced forgery or an official finding	Some indicators could fit deliberate imitation; no evidence of design has been found	"This was a false flag"; "planted to frame"; "staged"
No hostile act occurred (H3)	Named evidence classes examined, none hostile, and a positive cause identified	No hostile act found in the classes examined, which are listed; others remain	"No attack took place"; "ruled out" without the classes
Origin not established	None; this is the default state	Origin not established; here is what would change that	"Cannot be ruled out" with no evidence class attached

Source: Authors' proposal, Tier 3, applying Table 23. The example wording is illustrative and must be adapted to the facts of each case.

The third row is the hardest to word well. "There was no cyberattack" is a claim about everything. "No evidence of a cyber incident was found in these systems, after reviewing these evidence classes" is a claim about what was examined. The Iberian blackout committee's finding, as a press account reports it, is of the second kind: no indications or evidence of a cyber incident in the entities analyzed, after a review of IT systems, control-center protection systems, traces of malicious code and network traces, with vulnerabilities noted and no evidence that they were exploited (Section 8) [289]. The report itself was not read for this edition, so the example shows the form of the wording and does not vouch for the review behind it.

The fourth row deserves its own warning. "It cannot be ruled out" is true of nearly everything, so it carries almost no information, and it works as an insinuation when it follows a name or a motive. A statement that a hypothesis is open should say which class of evidence would test it and what is being done to collect that evidence.

The symmetry holds when the claim of staging comes from someone else. A statement that answers an adversary's pre-emptive "false flag" narrative says what evidence has been offered and what has been examined. It does not call the claim false unless an institution or court has examined it, which is the manual's reserved use of Refuted (Section 1).

D.3.5 Post-incident review

Every incident ends with a review that serves two purposes: to improve the response, and to start paying down the manual's largest debt, the lack of error rates. SP 800-61 Rev. 3 asks for an after-action report that documents the incident, the response and recovery actions taken, and the lessons learned (RC.RP-06), and it names lessons-learned meetings as a usual source of improvements (ID.IM-03) [403]. SP 800-86 observes that many forensic and incident response teams hold formal reviews after each major event and that such reviews usually lead to at least some changes (section 3.4) [405]. The annex adds three audits that look at what was said, what was set aside and what later proved right or wrong, and a fourth that checks the machinery (Table 75).

Table 75. Four audits after an incident

Audit	Question	Records used	Output
Statement audit	Did each public sentence stay within its gate and within what later proved true?	Statement register; gate log	Corrections issued; changes to wording templates
Hypothesis audit	Was any hypothesis set aside too early, or kept open too long?	Hypothesis log; discarded-hypotheses log	Revised tests and thresholds, with reasons
Call-and-outcome entry	What was called, with what confidence, and what resolved it?	Graded calls; resolution events when they occur	One register entry per call, updated over years
Process audit	Did the clocks, roles and handovers work as designed?	Intake sheet; custody log; contact log	Changes to plan, roles and training

Source: Authors' proposal, Tier 3, built on [403, 405]. The third audit is meant to yield error rates over time (Table 30, question 1).

The call-and-outcome register is the slow part. Resolution events arrive when a court rules, an authority reports or a seizure is announced, and in the manual's own cases the interval from event to first public attribution ran from at most 2 days to about 10 years, and to a final finding or admission from about seven weeks to about 7 years 9 months (Section 6.5). The register is therefore never closed. Each entry records the call, its date, its grades and its confidence, and a resolution field that stays open until something resolves it. Across incidents and organizations the entries could eventually yield the false-positive and false-negative rates that Table 30 names as the first open question (Section 11). The templates of D.4 are written so that organizations can pool entries without pooling the evidence (Tier 3).

The review also feeds three of the four governance indicators proposed in Section 8: the time from first signal to a joint review by human resources, security, compliance and legal functions, the share of incident reports that record a probability, a confidence level and at least one alternative hypothesis, and the share of reports that met the 24-hour, 72-hour and one-month windows. The hypothesis log and the statement register hold the raw material for all three.

Assessment

The register can show whether these procedures help only after enough calls have resolved, and the span from event to resolution means that will take years. Until then the annex claims auditability and no gain in accuracy (Tier 3).

D.4 Record templates

The procedures above rest on six records. Five are tables with fixed fields: the custody log (Table 60), the intake and clock sheet, the hypothesis log, the statement register and the call-and-outcome register. The sixth, the contact log, needs six fields only: time, counterpart, channel, what was shared, the agreement under which it was shared, and who approved it. Each record answers a question that an auditor would ask later. What did you hold? What did you believe? What did you say? Were you right? Blank copies of all six are published as CSV files with the open data of this edition under CC BY 4.0.

Records kept during an incident can be requested by a regulator, a court or an auditor, and a hypothesis log holds ideas that nobody has yet tested. SP 800-61 Rev. 3 asks that the confidentiality and integrity of response records be safeguarded and that only authorized personnel have access, with the incident lead often responsible for that (RS.AN-06) [403]. The format is meant to make the exposure safe. Entries are factual, dated, tied to item identifiers and labeled with their grade, so that a log read out of context says what the team knew and not what it feared. Legal should see the templates before an incident and not during one (Tier 3).

D.4.1 Intake and clock sheet

The intake sheet is the first record opened and the only one that stays live through all four gates. It fixes the two times that matter most, the time of first signal and the time of awareness, and it holds the notices due and sent. The NIS2 and GDPR clocks run from awareness (Table 66), so the reason for awareness matters as much as the time.

Table 76. Fields of the intake and clock sheet

Field	What to record	Basis
Incident ID	A unique number, never reused, carried by every record	Table 60
Time of first signal	UTC; who or what raised it	D.2.1, step 1
Time of awareness	UTC, and the reason the organization became aware	NIS2 Art. 23(4); GDPR Art. 33(1)
Reporting duties	Which notices apply, with the deadline time for each	Table 66
Notices sent	Recipient, time, content version and reference number	Table 66
Access conditions	Who could inspect the scene, logs or wreckage, and who could not	Section 8, step 2
Containment decisions	Action, time, reason, decider and effect on evidence	D.3.1
Roles assigned	A name for each role in Table 72	D.3.2
Gate status	G1 to G4, each open or closed, with date and reason	Table 73

Source: Authors' proposal, Tier 3, built on [303, 403, 407].

D.4.2 Hypothesis log

The hypothesis log is the internal record of what the team thinks and why. It carries H1, H2 and H3 of Table 23 from the first day, with the evidence that would count against each written in advance (Section 8, step 4). It also keeps the hypotheses the team has set aside, with the reasons, which is the discarded-hypotheses log of D.1.6. Nothing in it leaves the team.

Table 77. Fields of the hypothesis log

Field	What to record	Basis
Hypothesis	H1, H2, H3 or a local variant, in one sentence	Table 23
Claim and channel	Who asserted it, when and through which source	Section 8, step 1
Evidence for	Item identifiers, each with its tier and, for H1 and H2, its T, C and I grade	Section 8, step 5
Evidence against	Item identifiers, and the evidence that would count against it	Section 8, step 4
Alternatives	Plausible alternative explanations considered	SP 800-86, section 3.4
Status	Open, supported or set aside, with the date of each change	Section 8, step 7
Reason for change	Why the status changed and who agreed	Table 65

Field	What to record	Basis
Replication	Analyst B's result and every difference from Analyst A	Table 65
Next test	What would change the status, and who runs it	ICD 203, standards 2 to 4

Source: Authors' proposal, Tier 3, built on [281, 405]. Section references are to the seven-step procedure of this manual.

D.4.3 Statement register

The statement register is the one record whose output the public sees. It holds every statement as issued and ties each sentence to a gate and to evidence. An entry is made before release and never after.

Table 78. Fields of the statement register

Field	What to record	Basis
Statement	The exact words as issued, with channel and audience	Table 71
Time issued	UTC, and the time the draft was approved	Table 73
Gate	G1 to G4, and the record that opened it	Table 73
Five parts	Known; assessed with scale and confidence; unknown; what would change it; next update	Section 8
Approvals	Incident Lead, Legal and, at Gate 4, the Challenge Reviewer, each with a time	Table 72
Evidence references	The item identifiers behind each sentence	Table 60
Corrections	Later changes, each with date and reason	ICD 203, standard 7

Source: Authors' proposal, Tier 3, built on [281, 403]. The five-part form is the manual's own crisis-statement criterion.

D.4.4 Call-and-outcome register

This register is the slow one. Its job is to turn practice into data: each graded call is recorded when it is made and again when something resolves it. The events that resolve a call are of the kinds the manual finds behind its best-evidenced incidents, a court record, a seizure or a multi-state technical record (Section 6), and they can lag the call by weeks or by years.

Table 79. Fields of the call-and-outcome register

Field	What to record	Basis
Call	The hypothesis, the probability term, the confidence level and the T, C and I grades	Table 23
Date and gate	When the call was made and which gate it passed	Table 73
Scale used	ICD 203 table or the UK yardstick, named	ICD 203; PHIA yardstick
Basis	Evidence classes and tiers behind the call	Section 1
Resolution event	Court record, official report, seizure, admission, or none yet	Section 6
Outcome	Confirmed, contradicted, partly confirmed or unresolved	Table 30, question 1
Interval	Days from call to resolution, computed	Section 8

Source: Authors' proposal, Tier 3, built on [101, 281]. The register records calls and not the people who made them, because its purpose is to measure the method.

Pooling follows two rules. Entries are shared without evidence, victim data or staff names, so that no organization discloses more than the call, its grades and the outcome. And an entry is never deleted when a call proves wrong; it is closed with its outcome, since the wrong calls are the ones from which error rates are built (Tier 3).

D.5 Limits and reading route

This annex is a proposal, and its limits matter more than its procedures. A procedure that looks finished invites a trust that the evidence behind it does not yet support. Table 80 lists the limits the authors can name.

Table 80. Limits of Annex D and the route to improve each

Limit	What it means	Route to improvement
Untested	No gate, threshold or template has been run in a live incident or an exercise	Tabletop exercises with published after-action notes
No error rates	Indicator classes have no validated false-positive or false-negative rates (Table 30, question 1)	Pooled call-and-outcome registers over several years
No accuracy evidence	The one randomized test of a structured technique that was located found no gain in accuracy	Controlled tests of the replication and challenge steps
Indirect group-size basis	Marcoci et al. measured agreement on hypothetical reports and not accuracy of attribution	Studies on real attribution judgments
Uneven source depth	ISO/IEC 27037 read at informative level; Budapest party status unchecked; RFC 9162 is Experimental	Full texts, party checks and status tracking

Limit	What it means	Route to improvement
Law moves	NIS2 binds through national law; Spain's transposition and the 2020 guide are in flux	Re-check after transposition
Not legal advice	Preservation orders come from authorities; duties vary by country and sector	Counsel review before adoption
Narrow scope	Digital and open-source evidence only; physical scenes, human sources and classified material are outside	Separate protocols for those evidence types

Source: Authors' assessment, with Section 8 and Table 30 of this manual and [284, 406, 411, 412, 415].

Some readers will want three documents: an analytic manual, a forensic manual and a response manual. This edition keeps one, so that the evidence rules and the procedures that apply them cannot drift apart. The annex is written so that D.1 can be lifted out as a forensic manual and D.2 to D.4 as a response manual without changing Sections 1 to 11 (Tier 3). Table 81 shows the joins.

Table 81. Where Annex D meets the rest of the manual

Annex D part	Where it meets the manual	What it adds
D.1 Forensic protocols	Indicator classes in Table 22; the Berkeley Protocol as nearest standard (Section 8)	Handling that lets T, C and I be re-tested
D.2 Verification guide	Seven-step procedure (Section 8); NIS2 clocks in Table 25	Dated windows, evidence classes and statement ceilings
D.3 Response flows	Thresholds of Table 23; five-part statement (Section 8)	Roles, gates and symmetry wording
D.4 Templates	Governance indicators (Section 8)	The records from which the indicators are computed
Whole annex	Open questions 1 and 3 in Table 30 (Section 11)	A route toward error rates

Source: Authors' compilation from Sections 8 and 11 of this manual.

Key finding

The annex makes one claim that can be audited: for every sentence published about the cause of an incident or the actor behind it, the register can show the gate that was open and the record that opened it. Whether that discipline also makes the sentences more accurate is the question the call-and-outcome register exists to answer (Tier 3).

Readers with a specific job can start in different places. To decide what to say in public, read D.3.3, D.3.4, D.2.7 and D.2.8. For the first hours of an incident, go to D.2.1 and D.2.4. To build an evidence store, begin with D.1.1 and D.1.2. Read Table 80 before relying on any of it. Revisit the annex when Spain's NIS2 law takes effect, when NIST, ODNI or ISO revise the standards it leans on, and when the first call-and-outcome entries resolve.

Annex E. Spanish summary and bilingual glossary

Este anexo ofrece a los lectores hispanohablantes un resumen ejecutivo (E.1) y un glosario inglés-español (E.2). Es una traducción de apoyo: no añade datos ni conclusiones, y donde difiera del texto en inglés prevalece el texto en inglés. Las cifras del resumen se cotejaron una a una con la Table 1 y con las secciones citadas. «Operación negable» traduce deniable operation, es decir, una operación que el gobierno que la ordena puede desmentir con verosimilitud. Las calificaciones Documented, Contested, Alleged y Refuted aparecen como Documentado, Controvertido, Alegado y Refutado. Las tablas de este anexo siguen la numeración del manual.

Palabras clave: bandera falsa; guerra híbrida; guerra asimétrica; zona gris; atribución; manipulación de la información

E.1 Resumen ejecutivo

El Departamento de Seguridad e Inteligencia de ISDO preparó este manual para lectores que deben juzgar afirmaciones de que un incidente fue escenificado, disfrazado o atribuido falsamente a un tercero. Reúne tres cosas que suelen tratarse por separado: los conceptos de guerra asimétrica e híbrida, el registro jurídico y doctrinal, y la práctica de la atribución, con atención a Europa y a España. Su método es la clasificación de la evidencia en niveles. La doctrina declarada es el Nivel 1, la práctica verificada el Nivel 2 y la narrativa no confirmada el Nivel 3, y cada atribución se somete a las pruebas de trazabilidad técnica, convergencia de varias fuentes y evidencia de intención (Sección 1). El manual es analítico y defensivo, no toma partido político y no contiene ningún procedimiento para escenificar una operación ni para eludir la detección. La fecha de corte de la evidencia es el 11 de octubre de 2026.

Table 82. Cifras clave del manual (traducción de la Table 1)

Medida	Cifra	Base
Incidentes contemporáneos calificados en las tres pruebas	38	Registro de 2018 a 2026 a 11 de octubre de 2026 (Sección 6; Anexo B)
Incidentes que cumplen las tres pruebas, contando solo Y	1	Esquema norcoreano de trabajadores informáticos remotos (Sección 6)
Incidentes que cumplen al menos dos pruebas, contando solo Y	6	Cada uno llegó a un tribunal, a una incautación o a un registro de varios Estados (Sección 6)
Incidentes que no cumplen ninguna prueba, contando solo Y	20	De los 38; 12 cumplen una prueba y 5 cumplen dos (Table 50)

Medida	Cifra	Base
Mismos recuentos si una prueba cumplida en parte (P) cuenta como cumplida	17 cumplen las tres; 4 no cumplen ninguna	29 cumplen al menos dos; la regla está en el Anexo C.1.2
Del hecho a la primera atribución pública de un gobierno	De 2 días como máximo a unos 10 años	Crimea, Salisbury, Olympic Destroyer, TV5Monde (Sección 6.5)
Del hecho a la conclusión final o a la admisión	De unas 7 semanas a unos 7 años y 9 meses	Crimea, admisión en 2014; Salisbury, informe de la investigación en 2025 (Sección 6.5)
Afirmaciones calificadas frente a conclusiones institucionales	15: 8 Refutadas, 7 Alegadas	Conjunto elegido dentro de cuatro familias, no es una muestra (Sección 7)
Última sentencia localizada sobre el uso de banderas, insignias o uniformes enemigos	1947	Absolución de Skorzeny, 9 de septiembre de 1947 (Sección 4)
Incidentes FIMI contabilizados en 2025	540; 65 por ciento sin atribuir	4.º informe del SEAE, marzo de 2026 (Sección 8)
Total de sanciones híbridas de la UE en la renovación	80 personas, 20 entidades	Consejo, 8 de octubre de 2026; recuentos sin conciliar (Sección 9)
Contramedidas institucionales con evaluación de resultados	Ninguna en las fuentes leídas	Cinco medidas en la Table 24; otras dos tienen estudios de efecto (Sección 8)
Opciones de política valoradas por viabilidad	12: 2 altas, 6 medias, 4 bajas	Sección 10

Fuente: traducción de la Table 1; las cifras se cotejaron con ella una a una. FIMI es la manipulación e injerencia informativa extranjeras, y SEAE el Servicio Europeo de Acción Exterior. Los recuentos de incidentes descansan en las calificaciones de un solo lector, y cinco incidentes llevan la calificación R (fuente no recuperada), que podría trasladar hasta tres de ellos al grupo que cumple al menos dos pruebas (Anexo C.1.2).

Destacan ocho hallazgos.

- **La bandera falsa en sentido estricto es rara.** Con la regla de evidencia de la Sección 3, que exige evidencia de Nivel 2 sobre el diseño, los casos documentados son históricos: Mukden, Gleiwitz y, solo a través de fuentes secundarias, Mainila y el caso Lavon, cuya cadena de mando se discute. Olympic Destroyer y los correos de los Proud Boys de 2020 son los candidatos contemporáneos controvertidos. La mayoría de las operaciones examinadas recurre a intermediarios, contratistas e identidades simuladas y no culpa a ningún tercero; el registro calificado va de 2018 a 2026 (Sección 6).
- **La atribución llega por capas, y la evidencia sólida es la excepción.** Contando una prueba como cumplida solo cuando se califica Y, uno de los 38 incidentes contemporáneos cumple las tres pruebas, seis cumplen al menos dos y 20 no cumplen ninguna; si una prueba cumplida en parte también cuenta como

cumplida, 17 cumplen las tres (Anexo C.1.2). Una primera atribución pública de un gobierno llegó como máximo 2 días después del hecho en Crimea y 8 días después en Salisbury, pero 2 años y 8 meses después en Olympic Destroyer y unos 10 años después en TV5Monde. Una conclusión final o una admisión tardó unas siete semanas en Crimea y unos 7 años y 9 meses en Salisbury (Sección 6.5).

- **El derecho alcanza los elementos y no alcanza el diseño.** Cada régimen se aplica por encima de un umbral que una operación negable está construida para no cruzar, y todas las instituciones examinadas tratan como una decisión política soberana el hecho de nombrar a un Estado como responsable (Sección 4). Las sanciones de la UE contra personas eluden el umbral al incluir en listas a individuos y empresas.
- **El acceso, más que la retórica, separa las afirmaciones resueltas de las abiertas.** De quince afirmaciones elegidas dentro de cuatro familias, un conjunto que no es una muestra, ocho están Refutadas, cada una tras el examen de la afirmación concreta por una investigación, un tribunal o un órgano de inspección, y siete siguen Alegadas (Sección 7).
- **Los defensores tienen normas y pocas pruebas de resultados.** ICD 203 y el criterio de probabilidad del Reino Unido exigen hipótesis alternativas y un lenguaje calibrado, pero el único ensayo aleatorizado de técnicas estructuradas con analistas en activo que se localizó no halló mejora en la precisión (Sección 8). Solo dos medidas, la prohibición de RT y Sputnik y el pre-bunking, tienen estudios cuantificados de efecto, y ambos son de alcance limitado.
- **A las empresas no se les pide que nombren a los atacantes.** NIS2 fija plazos de 24 horas, 72 horas y un mes, y en el texto leído ninguna fase exige nombrar al actor. España reserva la atribución pública a un tercer Estado a su Consejo de Seguridad Nacional, y su proyecto de ley de NIS2 aún no es ley (Sección 8).
- **La negación plausible rinde en demora y en duda.** Nombrar a un actor tiene un coste para quien lo nombra, las listas híbridas de la UE son el único registro fechado de respuesta, y ninguna fuente muestra que alguna medida cambiara la conducta de un operador (Sección 9).
- **La evaluación va primero.** La Sección 10 valora doce opciones, dos de ellas con viabilidad alta. Cualquier opción puede entregarse como producto y ninguna tiene una métrica de resultados, de modo que el manual antepone el vocabulario compartido y los indicadores de resultados a los instrumentos nuevos.

Los analistas encontrarán en la Sección 8 un procedimiento de doble vertiente con umbrales, el registro completo en el Anexo B y el registro de verificación de fuentes en

el Anexo C. El Anexo D convierte el procedimiento en protocolos forenses, una guía de verificación para las primeras 24 horas, 72 horas y 30 días, y tres flujos de respuesta en paralelo; es una propuesta de los autores y no se ha probado (Nivel 3). La Sección 11 enumera cinco preguntas que la evidencia todavía no puede responder.

E.2 Glosario inglés-español

El glosario repite, en el orden alfabético inglés, los 58 términos de la Table 31 (Anexo A) con su equivalente español y una definición abreviada. Las definiciones resumen las inglesas y no las sustituyen. Las construcciones propias de los autores (Propuestas A a C, postura de atribución, procedimiento de doble vertiente, y puertas y regla de simetría del Anexo D) no tienen valor en el derecho ni en la doctrina.

Table 83. Glosario inglés-español de los 58 términos del manual

English term	Término en español	Definición breve en español
Access conditions	Condiciones de acceso	Quién pudo inspeccionar el lugar, los registros o los restos, y quién no.
Alleged (grade)	Alegado (calificación)	Afirmado pero no probado; marca una afirmación no contrastada y no significa que sea falsa.
Amplifier	Amplificador	Persona, medio o red que difunde el contenido o las afirmaciones de una operación.
Anchor (geolocation, chronolocation)	Ancla (geolocalización, cronolocalización)	Rasgo o registro verificable que fija dónde o cuándo se captó un elemento; el anexo pide dos.
Asymmetric warfare	Guerra asimétrica	Conflicto entre actores de fuerza desigual; la teoría pregunta por qué los más débiles a veces ganan.
Attribution	Atribución	Señalar al actor responsable de un acto, como juicio jurídico, técnico o político.
Attribution deception	Engaño de atribución	Propuesta C: todo esfuerzo deliberado por inducir a error a un público sobre quién es responsable.
Attribution posture	Postura de atribución	Variable de los autores, no probada: reconocida, gris, negada o falsificada, codificada con la regla de la Sección 5.8.
Black propaganda	Propaganda negra	Productos que parecen proceder de una fuente distinta de la verdadera.
Breakout Scale	Escala Breakout	Escala de seis categorías de la difusión de una operación, según su alcance entre plataformas y comunidades.
Call-and-outcome register	Registro de juicios y resultados	Registro de cada juicio emitido y, después, de su resultado, para estimar tasas de error.

English term	Término en español	Definición breve en español
Chain of custody	Cadena de custodia	Registro fechado de quién tuvo una prueba y qué se hizo con ella.
Challenge review	Revisión crítica	Intento independiente de refutar la hipótesis principal, documentado en un memorando.
Contested (grade)	Controvertido (calificación)	Bien probado, pero partes creíbles discuten la atribución, la intención o el alcance.
Countermeasures	Contramedidas	Jurídica: respuesta en otro caso ilícita a un hecho ilícito (artículo 49 de la CDI). Política: cualquier medida defensiva.
CSIRT	CSIRT (equipo de respuesta)	Equipo de respuesta a incidentes de seguridad informática; en NIS2, el órgano que recibe las notificaciones.
Cyber false attribution	Falsa atribución cibernética	Artefactos sembrados o prestados que hacen que una intrusión apunte a otro actor.
Deception	Engaño	Término doctrinal para las medidas deliberadas que inducen a error a quienes deciden.
Decision gate	Puerta de decisión	Punto de control que habilita una clase de afirmación solo cuando se cumplen sus condiciones de evidencia.
Deniability, plausible	Negación plausible	Conducir operaciones de modo que un gobierno pueda desmentirlas con verosimilitud (NSC 10/2, 1948).
Documented (grade)	Documentado (calificación)	Confirmado por admisión, sentencias, documentos desclasificados o varias investigaciones independientes.
Due diligence	Diligencia debida	Deber alegado de un Estado de no dejar que su territorio se use contra otros; es discutido.
Effective control	Control efectivo	Criterio de atribución que exige controlar la operación concreta; la CIJ rechazó el de control global.
Evidence tiers	Niveles de evidencia	Nivel 1 doctrina declarada, Nivel 2 práctica verificada, Nivel 3 narrativa no confirmada.
Evidentiary copy	Copia probatoria	Copia inalterada de un archivo fuente, con su hash, usada para el análisis y la réplica.
False colors	Pabellón falso	Uso por un buque del pabellón de otro Estado como estratagema; lícito hasta que empieza el ataque.
False flag (Proposal A)	Bandera falsa (Propuesta A)	Acto diseñado para que se atribuya la responsabilidad a otro actor concreto, en beneficio propio.
False flag, pretext subtype (Proposal B)	Bandera falsa, subtipo de pretexto (Propuesta B)	Operación de la Propuesta A cuya función principal es justificar una respuesta ya prevista.

English term	Término en español	Definición breve en español
False flag, registry sense	Bandera falsa, sentido registral	Buque que enarbola un pabellón al que no tiene derecho; no se traslada la culpa a nadie.
FIMI	FIMI	Manipulación e injerencia informativa extranjeras; el SEAE las contabiliza como incidentes.
Gray zone	Zona gris	Espacio operativo entre la paz y la guerra donde la coerción no cruza los umbrales de escalada.
Hash value	Valor hash	Resumen de longitud fija de los bytes de un archivo; cualquier cambio del archivo cambia el resumen.
Hybrid warfare	Guerra híbrida	Etiqueta discutida para modos de combate fusionados y gestión de umbrales; no hay definición acordada.
Hypothesis log	Registro de hipótesis	Registro fechado de cada hipótesis, sus pruebas y calificaciones, y por qué se mantuvo o se descartó.
ICD 203	ICD 203	Directiva de EE. UU. con nueve normas de oficio analítico, firmada el 2 de enero de 2015.
Kampala amendments	Enmiendas de Kampala	Enmiendas de 2010 que dan a la CPI competencia sobre la agresión; activas desde 2018.
Liar's dividend	Dividendo del mentiroso	Ganancia de llamar falso o escenificado, sin base, a un hecho real (Chesney y Citron).
Masked identity	Identidad enmascarada	Personaje prestado o inventado que oculta al actor real.
Metadata state	Estado de los metadatos	Una de cinco condiciones de los metadatos de un archivo, de original a no resuelto.
Misdirection (post-attack)	Desvío (tras el ataque)	Funcionarios que apartan una investigación de los verdaderos autores después de un atentado.
NIS2	NIS2	Directiva (UE) 2022/2555; plazos de notificación de 24 horas, 72 horas y un mes.
Non-intervention	No intervención	Regla contra coaccionar las decisiones libres de otro Estado; la mera persuasión no basta.
Order of volatility	Orden de volatilidad	Secuencia para recoger pruebas, de los datos más perecederos a los más duraderos.
Passive collection	Recopilación pasiva	Reunir solo lo que tiene el defensor o publican terceros, sin contacto con sistemas sospechosos.
Perfidy	Perfidia	Simular un estatus protegido para traicionar la confianza del adversario; prohibida por el artículo 37 del PA I.
Person-based sanctions	Sanciones a personas	Listas de la UE de personas y entidades por lo que hicieron, sin ninguna conclusión sobre un Estado.
PHIA yardstick	Escala PHIA	Escala de probabilidad británica del Jefe Profesional de Evaluación de Inteligencia, que une palabras y rangos.

English term	Término en español	Definición breve en español
Pre-bunking	Pre-bunking (advertencia previa)	Avisar a las personas de las técnicas de manipulación antes de la exposición, con vídeos o juegos breves.
Pretext	Pretexto	Incidente, real o escenificado, usado para justificar una respuesta ya prevista.
Refuted (grade)	Refutado (calificación)	Contradicho por una institución o un tribunal que examinó la afirmación concreta.
Ruse of war	Estratagema de guerra	Engaño lícito como el camuflaje, los señuelos y la desinformación; no está prohibido (artículo 37 del PA I).
State responsibility	Responsabilidad del Estado	Derecho que atribuye una conducta a un Estado y fija sus consecuencias (artículos de la CDI, 2001).
Statement register	Registro de declaraciones	Registro de cada declaración pública sobre un incidente, con su redacción, su techo y quien la aprueba.
Stay-behind network (Gladio)	Red stay-behind (Gladio)	Red secreta vinculada a la OTAN y oculta a los parlamentos; existencia documentada, papel en atentados no probado.
Symmetry rule	Regla de simetría	Afirmar una escenificación exige la misma evidencia que afirmar que lo hizo un actor nombrado.
T/C/I tests	Pruebas T/C/I	Trazabilidad técnica, convergencia de varias fuentes, evidencia de intención; cada una se califica Y, P, N, U o R.
Three Warfares	Tres Guerras	Doctrina del EPL de guerra de opinión pública, psicológica y jurídica; texto no leído.
Two-sided procedure	Procedimiento de doble vertiente	Método de los autores que contrasta tres hipótesis: actuó la parte nombrada, escenificación o disfraz, ningún acto hostil.

Fuente: traducción abreviada de la Table 31 (Anexo A); ante cualquier discrepancia prevalece el texto en inglés. PA I es el Protocolo Adicional I de 1977 a los Convenios de Ginebra; CDI, la Comisión de Derecho Internacional; CIJ, la Corte Internacional de Justicia; CPI, la Corte Penal Internacional; EPL, el Ejército Popular de Liberación de China; SEAE, el Servicio Europeo de Acción Exterior.