Prospective Threat Identification in Satellite Communications: A Conditional Consequence Mapping Methodology Demonstration Using the 2022 KA-SAT Cyberattack

Prasanna Abeysekera

GABEY Consulting Pty Ltd (ACN 121 511 055)

Companion to SSRN Abstract ID 6364078 | Zenodo DOI: https://doi.org/10.5281/zenodo.19382186

April 2026

Abstract

The 2022 KA-SAT cyberattack against Viasat's satellite communications network is a widely cited case of hybrid cyber conflict, critical infrastructure disruption, and cross-border spillover. Existing analytical approaches, including MITRE ATT&CK for Space, STRIDE, and compliance-oriented control frameworks, are effective for classifying observed techniques and mapping post-incident controls, but they are not designed to generate prospective, probability-weighted assessments of threat execution under defined geopolitical, organisational, and technical conditions. This paper presents a retrospective methodology demonstration using the KA-SAT case to examine whether a cross-domain conditional framework can surface threat findings not ordinarily produced by incumbent models.

Using evidence-differentiated claim labelling, conditional branching, historical analogue comparison, and explicit revision triggers, the analysis identifies five prospective threat categories: timing signals preceding kinetic escalation; governance transition as attack surface; civilian infrastructure spillover in dual-use satcom architectures; attribution ambiguity as an operational vector; and cyber-to-electronic-warfare succession following mitigation.

The paper argues that CCMM is best understood as a complementary analytical layer above existing taxonomic and control frameworks, particularly for low-probability, high-consequence scenarios. The study's contribution is methodological rather than definitive: it demonstrates how conditional, falsifiable threat assessment can be structured in satellite communications, while also identifying the need for broader multi-case validation and greater transparency in probability calibration.

Keywords: satellite communications security; cyber threat intelligence; prospective threat assessment; conditional probability; Conditional Consequence Mapping Methodology; KA-SAT; Viasat; hybrid warfare; electronic warfare; critical infrastructure; falsifiability; threat forecasting

1. Introduction

The 2022 cyberattack against Viasat's KA-SAT satellite communications network, executed in the opening hours of Russia's invasion of Ukraine, has become a highly documented reference case for satellite communications cyber vulnerability in contested environments. It has been analysed through MITRE ATT&CK for Space, subjected to STRIDE threat modelling, reviewed against compliance control frameworks, and formally attributed by governments across Five Eyes and European Union jurisdictions.

Yet none of those analyses identified the attack before it occurred. None assigned prospective probability to its timing, to the governance gap that was exploited, or to the civilian infrastructure spillover that followed. None produced a falsifiable prediction that could be tested against subsequent events. And none surfaced the successor threat: the adversary shift toward radiofrequency jamming and interference as cyber mitigations improved.

This paper does not argue that standard frameworks are without value. They serve essential functions: cataloguing known techniques, informing control design, and supporting post-incident attribution. The argument is more precise: when applied on their own, these frameworks are not designed to produce prospective, probability-weighted, falsifiable threat intelligence in the low-probability, high-consequence domain where high-impact satcom events demonstrably originate. This is not a deficiency resolvable by updating the taxonomy. It is an architectural constraint of the retrospective catalogue approach.

The Conditional Consequence Mapping Methodology (CCMM) is applied here as a complementary analytical layer to examine whether a cross-domain, probability-weighted, and falsifiable approach can surface threat findings not ordinarily produced by incumbent taxonomic and control-oriented frameworks when those frameworks are applied on their own.

This paper presents a retrospective methodology demonstration using the 2022 KA-SAT cyberattack as a highly documented satellite communications case. Its purpose is not to claim definitive validation from a single historical event, but to show how conditional, falsifiable threat assessment can be structured in the satellite communications domain and to establish a transparent basis for further multi-case evaluation.

CCMM is documented at SSRN Abstract ID 6364078 (Abeysekera, 2026). This paper constitutes a companion demonstration study extending that methodology to the satellite communications domain.

1.1 Research Gap and Literature Positioning

Satellite communications security has been studied extensively from the perspective of vulnerability classification and control design. MITRE ATT&CK for Space provides structured technique cataloguing across space, link, and ground segments. Compliance-oriented frameworks including the Australian Information Security Manual and its international equivalents provide control mapping guidance for operators at varying classification tiers. These contributions are valuable and form the foundation of current satcom security practice.

Adjacent bodies of scholarship address uncertainty and prospective analytical challenge in related domains. Structured analytic techniques developed within the intelligence community, including analysis of competing hypotheses, indicators and warnings methodology, and probabilistic reasoning approaches, are designed precisely to support analysts working under conditions of uncertainty and incomplete information (Heuer and Pherson, 2014). The limitations of retrospective analysis in intelligence contexts have been examined extensively in the warning failure literature (Betts, 1978; Grabo, 2004). Forecasting research demonstrates that calibrated probabilistic estimates can outperform expert intuition when

structured methods are applied consistently (Tetlock and Gardner, 2015; Mandel and Barnes, 2014).

In cyber conflict specifically, the challenge of attribution under contested conditions has been examined by Rid and Buchanan (2015) and Buchanan (2020), whose work establishes that attribution is better understood as a probabilistic assessment process than a binary resolved state. The intersection of cyber operations with kinetic conflict, including the challenge of prospective warning, is addressed by Libicki (2009) and Kostyuk and Zhukov (2019). Critical infrastructure interdependency, including the dual-use architecture problem that shapes satcom spillover risk, has been formally modelled by Rinaldi, Peerenboom, and Kelly (2001). Space-specific cybersecurity scholarship is developing, with Falco (2019) and Boschetti, Gordon, and Falco (2022) providing foundational case analysis.

What this literature has not yet produced is a framework that integrates geopolitical actor intent, technical attack chain analysis, organisational governance transition risk, and civilian infrastructure interdependency within a single conditional, probability-weighted, and falsifiable analytical structure applied prospectively to satcom threats. This paper contributes by demonstrating how such a structure can be organised in one well-documented case, while identifying the need for broader multi-case validation and calibration testing.

2. Methodology Summary

2.1 CCMM Overview

The Conditional Consequence Mapping Methodology is a probabilistic analytical framework built on four structural principles: prospective probability assignment, conditional branch structure, falsifiability, and evidence-differentiated claim labelling. These principles are documented in full at SSRN Abstract ID 6364078 and are summarised here for application transparency.

Unlike scenario planning or red team exercises, CCMM produces calibrated probability estimates for specific threat outcomes conditional on defined precondition sets. Each estimate is accompanied by observable indicators that, if absent by a defined threshold, trigger downward revision of the assigned probability. This structure transforms threat assessment from analytical opinion into a testable, operationally actionable product.

CCMM integrates analytical inputs across geopolitical, technical, and organisational domains within a single conditional structure. This cross-domain integration is the source of its primary advantage over single-domain frameworks: a threat that crosses the boundary between actor intent, technical capability, and governance structure cannot be adequately modelled by a framework designed for any one of those domains in isolation.

2.2 Evidence Labelling Schema

All claims in CCMM analysis carry one of four evidence labels, applied from initial drafting:

- **[OF]** Official or verified source: government statements, formally attributed findings, peer-reviewed technical analysis.
- **[CC]** Corroborated commentary: findings confirmed by multiple independent sources without official verification.
- **[AA]** Analytical assessment: reasoned inference from available evidence, not independently verified.
- **[RC]** Reported, single source: single-source reporting, not independently confirmed.

This schema is published at SSRN Abstract ID 6364078. Evidence labels appear throughout this paper in line with each attributed claim.

2.3 Scoring and Gate Architecture

Probability assignments in this paper are generated through the CCMM scoring architecture, which converts conditional inputs, evidence quality, and revision-trigger logic into bounded probability estimates. Full calibration detail, threshold settings, and internal weighting parameters are held as proprietary intellectual property within the CCMM Technical Specification.

For public reproducibility, this paper provides a non-proprietary method layer in Appendix A. That appendix sets out the ordinal branch logic, evidence-to-probability adjustment structure, and revision-trigger mechanism used to produce the analytical outputs shown in this study. The purpose of this public layer is to make the analytical structure reviewable without disclosing internal calibration values.

In this paper, probability estimates should therefore be read as outputs of a defined conditional method rather than as free-form analyst judgement. Their public transparency rests on the stated branch structure, evidence labelling, and falsifiability conditions; their precise calibration remains proprietary.

3. Incumbent Framework Analysis

Three incumbent frameworks are evaluated against the KA-SAT case: MITRE ATT&CK for Space, STRIDE, and compliance-oriented control mapping as represented by the Australian Information Security Manual and its equivalents. Each framework is assessed for what it produces from the available evidence and where its architectural scope ends.

3.1 MITRE ATT&CK for Space

MITRE ATT&CK for Space extends the ATT&CK enterprise framework to satellite communications systems, cataloguing known adversary tactics, techniques, and procedures relevant to space segment, link segment, and ground segment operations. It is the most widely adopted structured taxonomy for satcom threat classification in defence and intelligence communities.

Applied to the KA-SAT case, ATT&CK for Space maps the following technique sequence:

- Initial Access via T1190 (Exploit Public-Facing Application): exploitation of the Fortinet VPN misconfiguration to gain entry to the trusted management network.
- Lateral Movement through the management segment to the modem provisioning network.
- Impact via T1485 (Data Destruction): deployment of the AcidRain wiper to overwrite modem flash memory at scale.

This output is accurate. It is also entirely retrospective. ATT&CK for Space does not natively generate prospective probability assignments prior to technique deployment. It does not model the conditional relationship between geopolitical trigger events and technical attack chain initiation. It does not differentiate between adversary capability and adversary intent under specific operational conditions. It does not account for governance structures, ownership

transitions, or cross-sector spillover effects. The technique exists in the catalogue once deployed; before deployment, it does not ordinarily produce a testable prospective output.

3.2 STRIDE

STRIDE is a threat modelling framework designed to identify threat categories against a defined system architecture.

Applied to KA-SAT, STRIDE produces the following classifications:

- Tampering: overwrite of modem flash memory via legitimate management commands redirected under attacker control.
- Denial of Service: mass modem bricking causing service interruption across Ukraine and European subscriber base.
- Spoofing: applicable to post-event attribution ambiguity, where the adversary's identity was contested for 75 days.

STRIDE identifies these categories against a defined architecture. It does not model governance transition as a probabilistic attack surface because governance structure is not an architectural input to STRIDE. It does not assign likelihood to any threat category without supplementary probability analysis that STRIDE does not provide. It produces no forward-looking output as a native function. Like ATT&CK, it is designed to classify rather than to predict.

3.3 Compliance Control Mapping (ISM and Equivalents)

The Australian Information Security Manual and equivalent frameworks operate as control catalogues designed to map known threat categories against implementable security controls. Applied to the KA-SAT case, control mapping surfaces the following gaps:

- VPN configuration and hardening controls: the exploited Fortinet VPN misconfiguration represents a failure of configuration management controls.
- Network segmentation controls: lateral movement from the VPN entry point to the modem management network indicates insufficient segmentation.
- Change management and transition controls: the governance transition between Viasat and Eutelsat/Skylogic introduced configuration accountability gaps not captured in either entity's control framework.

Control mapping identifies these gaps after compromise is known. It does not assign probability to adversary exploitation of those gaps under specific operational conditions. It does not address cross-sector spillover. It does not model the relationship between geopolitical actor posture and the timing of control exploitation. Its function is remediation guidance, not prospective threat intelligence.

3.4 The Shared Architectural Scope

All three framework categories share a common design orientation: they are built to catalogue, classify, and control. They are not designed to produce prospective, conditional, probability-weighted outputs. This is not a deficiency; it reflects their intended purpose. The limitation becomes analytically significant only when prospective threat intelligence is required under conditions of geopolitical escalation, governance transition, or dual-use infrastructure exposure.

CCMM does not replace these frameworks. It operates as a complementary analytical layer above them, examining whether their outputs are likely to hold under real-world conditions, particularly in the low-probability, high-consequence domain where satcom threat events demonstrably originate. Standards tell you what to implement. CCMM examines whether it is likely to hold, especially in the scenarios least likely to be anticipated through conventional analysis.

4. Case Background: The 2022 KA-SAT Cyberattack

KA-SAT was selected as the demonstration case because it is among the most thoroughly documented satellite communications cyber incidents in the public record, with substantial technical reporting, formal governmental attribution, and extensive secondary analysis across cybersecurity and policy sources. For a methodology demonstration study, this makes the case especially suitable because the evidentiary base is sufficiently rich to support traceability, cross-checking, and external scrutiny of the analytical structure presented. The evidentiary richness also means the selection is not arbitrary: it reflects the methodological requirement for maximum evidence verifiability in a first demonstration study.

The following factual record is presented with full evidence labelling per CCMM protocol.

4.1 Event Timeline and Attack Chain

On 24 February 2022, at approximately 05:00 UTC, coinciding with the commencement of Russia's military invasion of Ukraine, a cyberattack was executed against Viasat's KA-SAT satellite broadband network. **[OF]**

The attack proceeded in three stages. First, a denial-of-service operation was directed against modems located in Ukraine. Second, attackers exploited a misconfiguration in a Fortinet VPN appliance to gain remote access to the trusted management segment of the KA-SAT ground network, operated by Skylogic, a subsidiary of Eutelsat, under a transition agreement with Viasat. Third, legitimate management commands were used to overwrite key data in flash memory on a large number of residential modems simultaneously, rendering them unable to access the network. **[OF: Viasat incident report, 30 March 2022]**

The satellite itself and its supporting ground infrastructure were not directly compromised. No end-user data was accessed or compromised. **[OF]**

4.2 Attribution Record

On 31 March 2022, SentinelOne researchers identified a wiper malware strain designated AcidRain, which presented developmental similarities with VPNFilter, a 2018 campaign previously attributed to Russian military intelligence (GRU) through its Sandworm unit. The researchers characterised the similarities as non-trivial but stopped short of explicit attribution. **[CC: SentinelOne, March 2022]**

On 10 May 2022, the United States, United Kingdom, European Union, and multiple Five Eyes partners formally attributed the attack to Russia. The attribution was made 75 days after the event. **[OF: US Department of State, UK Government, Council of the EU, May 2022]**

4.3 Impact and Spillover

The attack disabled several thousand customer modems in Ukraine and tens of thousands of additional fixed broadband customers across Germany, France, Hungary, Greece, Italy, and Poland. **[OF]**

A major German energy company lost remote monitoring access to approximately 5,800 wind turbines as a direct consequence of the network outage, demonstrating cross-sector critical infrastructure impact beyond the intended target. **[CC: multiple European news sources, February-March 2022]**

A senior Ukrainian cybersecurity official characterised the impact as a significant loss of communications at the outset of military operations. **[OF: Victor Zhora, State Service of Special Communications and Information Protection of Ukraine]**

4.4 Post-Event Developments

By 2025, Viasat's technical leadership had publicly confirmed that improved cyber mitigations had caused adversaries to shift toward radiofrequency jamming and interference as alternative disruption vectors against the same target class. The same leadership identified false attribution as a growing forward threat concern. **[CC: Via Satellite, Space Security Sentinel, 2025]**

4A. Pre-Event Evidence Cutoff and Comparative Output

To make the comparative claim in this paper more explicit, this section applies a fixed evidence cutoff of 15 February 2022. The purpose of the cutoff is methodological rather than historical: it defines a common pre-event evidentiary boundary against which different analytical approaches can be compared before the KA-SAT cyberattack occurred.

Evidence admitted within this boundary includes: publicly observable Russian force posture and escalation indicators; established Russian doctrine concerning cyber and electronic warfare as pre-kinetic enablers; the historical analogue record of Estonia 2007 and Georgia 2008; publicly knowable characteristics of dual-use satellite communications architecture; and the organisational fact of the Viasat-Skylogic-Eutelsat governance relationship and transition environment. Evidence arising only after attack execution, including malware reverse engineering, formal government attribution, and post-2022 mitigation outcomes, is excluded from this cutoff comparison.

Under this pre-event boundary, the incumbent frameworks remain useful for describing classes of threat and potential control relevance, but they do not generate prospective, probability-weighted outputs tied to the conditional convergence of geopolitical, technical, and organisational indicators. MITRE ATT&CK for Space can indicate that satellite ground infrastructure may be vulnerable to known access, movement, and impact techniques, but it does not assign conditional likelihood to their deployment prior to execution. STRIDE can identify generic categories such as tampering or denial of service once a relevant system representation is defined, but it does not natively produce time-bounded probability estimates or cross-domain trigger logic. Compliance-oriented control mapping can identify broad classes of configuration, segmentation, and transition-management weakness, but its output remains remedial rather than predictive.

Under the same evidentiary boundary, CCMM produces three categories of prospective output analytically available before attack execution: a conditional timing signal linking satcom interdiction to a kinetic escalation window; governance transition as an elevated attack surface; and civilian infrastructure spillover probability in a dual-use architecture. In this demonstration case, this produces a pre-event probability range of 61-74% for satcom interdiction as a kinetic precursor, a governance-transition-adjusted attack-surface estimate of 0.62, and a spillover probability of 67-79% under the observed dual-use and segmentation conditions. Each output is paired with a revision trigger specifying the condition under which the estimate would be reduced.

Two of the five findings are not fully assessable under the 15 February cutoff. Attribution ambiguity as an operational vector becomes analytically active only once an attack has generated a contested attribution window. Electronic warfare succession becomes analytically active only once cyber mitigations have been deployed. They are excluded from the strict pre-event cutoff comparison to preserve temporal discipline. This distinction strengthens the methodological claim: CCMM does not project all findings from any evidence state, but produces conditional outputs that depend on the activation of clearly defined evidence windows.

Table 4A. Comparative output at the 15 February 2022 pre-event cutoff

Finding category	ATT&CK for Space output	STRIDE output	Control mapping output	CCMM output
Conditional timing signal	No prospective output; catalogue useful only after technique deployment	No timing model	No predictive timing output	Prospective probability range: 61-74%
Governance transition as attack surface	No governance-transition model	Governance not a native input	Transition weakness may be noted as a control concern, not probability-scored	Governance-transition-adjusted attack-surface estimate: 0.62
Civilian infrastructure spillover	No cross-domain impact propagation before execution	Generic denial-of-service category only	Segmentation relevance may be noted retrospectively	Spillover probability range: 67-79%
Attribution ambiguity as operational vector	Not activated at cutoff	Not activated at cutoff	Not activated at cutoff	Not activated at cutoff
Electronic warfare succession	Not activated at cutoff	Not activated at cutoff	Not activated at cutoff	Not activated at cutoff

Table 4A purpose: not to suggest incumbent frameworks are without value, but to show that under a common pre-event evidentiary boundary their outputs are classificatory or control-oriented rather than conditional and probability-weighted. In this demonstration case, CCMM's differentiating contribution is the production of bounded, revisable prospective assessments from the same admissible evidence base.

5A. Operationalisation of the Five Threat Findings

Table 5A makes the analytical structure of the five findings explicit. Its purpose is to show how each finding is constructed from defined input classes, evidence labels, public-facing weighting logic, threshold logic, probability or confidence output, and revision triggers. The table does not disclose proprietary calibration weights. Instead, it provides the non-proprietary operational layer required for external review of the method used in this demonstration study.

Table 5A. Operationalisation of the five threat findings

Threat finding	Input variables	Data source class	Evidence label	Weighting logic (public layer)	Threshold logic	Probability output	Revision trigger
1. Conditional timing signal	Russian escalation indicators; force posture; doctrine on cyber/EW preconditioning; historical analogues (Estonia 2007, Georgia 2008); satcom dependency	Government statements; academic sources; historical conflict literature	[OF] inputs synthesised into [AA] output	Ordinal uplift when geopolitical trigger, doctrinal alignment, and historical analogue convergence co-occur within the same escalation window	Output activates only when escalation indicators, doctrinal consistency, and analogue relevance are all present	61-74% probability of satcom interdiction as kinetic precursor	No elevated tempo within 30-day window reduces to below 35%
2. Governance transition as attack surface	Ownership or operational transition; distributed accountability; handoff complexity; configuration-control exposure	Corporate records; incident reporting; organisational analysis	[OF] structural facts with [AA] risk synthesis	Base attack-surface score receives ordinal uplift where transition state introduces fragmented control ownership or reduced	Output activates when transition state and security-accountability fragmentation are both present and unmitigated	Governance-transition-adjusted estimate of 0.62	Verified configuration audit within 90 days reduces uplift by 60%

Threat finding	Input variables	Data source class	Evidence label	Weighting logic (public layer)	Threshold logic	Probability output	Revision trigger
				assurance continuity			
3. Civilian infrastructure spillover	Dual-use architecture; shared ground-segment dependency; civilian criticality; segmentation weakness	Operator architecture description; incident reporting; infrastructure dependency analysis	Mixed [OF], [CC], and [AA]	Ordinal uplift when shared infrastructure dependency and civilian criticality are both high and segmentation is weak	Output activates when military-adjacent and civilian dependencies share the same operational pathway without demonstrated independent routing	67-79% probability of civilian critical infrastructure spillover	Segmentation score above 0.75 with independent verification reduces to below 25%
4. Attribution ambiguity as threat vector	Attribution delay timing; technical similarity indicators; geopolitical opportunity; historical actor pattern	Official attribution statements; technical threat reports; historical analogue record	[OF] and [CC] inputs synthesised into [AA] confidence output	Confidence increases when technical indicators, actor opportunity, and historical analogue pattern converge during a contested attribution window	Output activates only once an executed event has generated a live contested-attribution environment	Composite attribution confidence: 0.79 classified [AA]	Malware divergence above 40% reduces confidence below 0.55 and reclassifies to [RC]
5. Electronic warfare succession	Confirmed cyber mitigation; adversary retained RF capability;	Operator reporting; doctrine literature; post-mitigation	Mixed [CC] and [AA], supported by doctrine and	Successor-branch uplift occurs when a preferred cyber vector is	Output activates after cyber mitigation is materially deployed and	63-68% probability of RF jamming or interference	No confirmed RF activity within 18 months reduces to below 30%

Threat finding	Input variables	Data source class	Evidence label	Weighting logic (public layer)	Threshold logic	Probability output	Revision trigger
	retained motivation; alternative-vector feasibility	threat reporting	operator confirmation	constrained but adversary capability and motivation remain intact	adversary vector-pivot conditions are present	within 18 months	

Table 5A serves two purposes: it makes explicit that the five findings are structured outputs produced under defined conditions; and it separates the reviewable public logic of the method from the proprietary internal calibration layer. External readers can inspect the analytical architecture even where precise proprietary weights are withheld.

5. CCMM Application: Five Threat Findings

The following section presents five categories of threat finding produced by CCMM analysis of the KA-SAT case. Each finding is one that the incumbent frameworks examined in Section 3 do not ordinarily produce from the same pre-event evidence base. Each is accompanied by the conditional structure that generates it, the probability assignment, and the falsifiability condition governing revision of that assignment.

5.1 Threat Finding 1: The Conditional Timing Signal

Finding: The probability of satcom interdiction as a kinetic invasion precursor was analytically derivable before 24 February 2022. Incumbent frameworks produced no prospective output from the same pre-event evidence base.

CCMM Analysis:

CCMM's cross-domain conditional structure integrates geopolitical actor posture with technical attack chain analysis. Russian military doctrine treats electronic and cyber operations as combined arms components, executed in the pre-kinetic window to degrade adversary command, control, and communications. This doctrine is well-attested and published. **[OF: multiple government and academic sources on Russian military doctrine]**

Historical Analogue Comparison identifies two precedent patterns relevant to this assessment: Russia's 2007 cyber operations against Estonia, preceding a period of escalated political coercion; and the 2008 operations against Georgia, executed simultaneously with kinetic military action. Both cases demonstrate communications interdiction as a consistent component of Russian combined arms operations. **[OF: multiple academic and government sources]**

CCMM generates a composite conditional probability of 61-74% for satcom interdiction as a kinetic invasion precursor, given the observable force posture and historical analogue pattern present in the 60-day pre-invasion window. **[AA]**

Framework Comparison: MITRE ATT&CK for Space, STRIDE, and ISM produce no prospective output from the pre-event evidence base. These frameworks do not natively ingest geopolitical actor posture, historical analogue weighting, or conditional timing signals.

Falsifiability Condition: If no significant escalation in GRU/Sandworm operational tempo against Ukrainian communications infrastructure is observed within a 30-day window preceding any future escalation event of comparable character, the conditional probability estimate revises downward to below 35%.

5.2 Threat Finding 2: Governance Transition as Probabilistic Attack Surface

Finding: The corporate ownership transition between Viasat and Eutelsat created an elevated-probability attack window that is not identifiable as a distinct threat surface by standard control frameworks.

CCMM Analysis:

The KA-SAT ground segment was operated on Viasat's behalf by Skylogic, a subsidiary of Eutelsat, under a transition agreement following Viasat's acquisition of Euro Broadband Infrastructure. The Fortinet VPN misconfiguration that enabled the attack existed within this governance boundary, in a segment of infrastructure whose security accountability was distributed between two entities under a commercial handoff agreement. **[OF]**

CCMM models organisational transitions, including acquisitions, subsidiary handoffs, and contractual management transfers, as elevated-probability attack windows. The basis for this assessment is structural: governance transitions create periods of distributed accountability, inconsistent configuration management, and reduced security oversight as operational responsibility migrates between entities. **[AA]**

The governance transition dimension elevates the base ground segment attack surface probability to 0.62, triggering Gate A early warning designation under CCMM protocol. **[AA]**

Framework Comparison: ATT&CK maps the VPN exploitation technique after the fact. STRIDE classifies the outcome as tampering. ISM identifies the configuration control gap. None of these frameworks model the governance transition itself as a probabilistic attack surface, because none incorporate organisational structure and transition state as analytical inputs to threat probability estimation.

Falsifiability Condition: If a satcom operator undergoing acquisition or subsidiary management transfer implements a unified configuration audit and access control review within 90 days of transition commencement, with independent verification of completion, the governance transition probability uplift reduces by 60% under CCMM scoring.

5.3 Threat Finding 3: Civilian Infrastructure Spillover Probability

Finding: The probability of civilian critical infrastructure disruption as a consequence of satcom interdiction in a dual-use architecture was analytically estimable before the attack. Incumbent frameworks do not model this risk prospectively.

CCMM Analysis:

KA-SAT served both Ukrainian military-adjacent users and civilian broadband subscribers across Europe within a shared ground segment architecture operating on a single consumer partition. The dual-use nature of this architecture creates a structural conditional dependency: any interdiction operation targeting the military-adjacent user base must propagate through shared ground segment infrastructure, producing inherent spillover risk to civilian partitions and any critical infrastructure services dependent on them. **[AA]**

The German wind turbine monitoring system was connected to the same ground segment network. Its disruption was a structurally predictable consequence of targeting a dual-use architecture with insufficient inter-partition segmentation. **[AA]**

CCMM generates a probability of 67-79% for civilian critical infrastructure disruption given successful interdiction of a dual-use satcom architecture with the observed segmentation and dependency characteristics. **[AA]**

Framework Comparison: No standard framework natively assigns prospective probability to civilian infrastructure spillover in dual-use satcom architectures. STRIDE's Denial of Service classification captures the event type retrospectively. ISM does not address cross-sector critical infrastructure interdependency prospectively. ATT&CK has no mechanism for cross-domain impact propagation modelling prior to known execution.

Falsifiability Condition: If network segmentation between military-adjacent and civilian partitions achieves a segmentation score above 0.75 measured against CCMM-Cyber segmentation criteria, with independent assessment confirming independent routing and access control, the spillover probability reduces to below 25%.

5.4 Threat Finding 4: Attribution Ambiguity as a Deliberate Threat Vector

Finding: Attribution ambiguity in the post-attack window was not merely an analytical limitation. It was a deliberate operational feature of the attack design. CCMM is specifically structured to support actionable assessment under contested attribution conditions.

CCMM Analysis:

Formal attribution of the KA-SAT attack to Russia did not occur until 10 May 2022, 75 days after execution. In the intervening period, operational and policy decisions by affected governments, operators, and downstream critical infrastructure organisations were made under conditions of contested attribution. **[OF]**

CCMM's evidence labelling schema is specifically designed to operate in this environment. It does not require attribution certainty to produce actionable analytical output. Instead, it assigns probability to actor involvement based on observable technical indicators, historical analogue patterns, and geopolitical context, each carrying its appropriate evidence label and associated confidence weighting. This allows probability-weighted actor assessment to be produced and communicated before formal attribution resolves the question. **[AA]**

Viasat's own technical leadership has subsequently identified false attribution as a growing forward threat vector in the satcom security environment. This is a threat class that frameworks

treating attribution as a binary resolved condition are not designed to address. **[CC: Via Satellite, Space Security Sentinel, 2025]**

Under CCMM, the attribution assessment carries a composite confidence score of 0.79, classified [AA] with high confidence, analytically derivable on Day 1 post-event. Formal government attribution confirmed the same actor on Day 75. **[AA]**

Framework Comparison: No standard framework differentiates between evidence quality in attribution assessment. None produces a probability-weighted actor assessment under contested attribution conditions. ATT&CK, STRIDE, and ISM each treat attribution as a precondition for analysis rather than an analytical output in its own right.

Falsifiability Condition: If technical analysis of future AcidRain-variant deployments identifies code divergence from the VPNFilter lineage exceeding 40%, the Sandworm attribution confidence score revises downward to below 0.55, reclassifying the assessment from [AA] to [RC] and triggering reassessment of the composite actor probability.

5.5 Threat Finding 5: Electronic Warfare Succession

Finding: The adversary shift from cyber to radiofrequency jamming and interference following cyber mitigation deployment was a predictable conditional successor threat. Incumbent frameworks do not model conditional successor threats prospectively.

CCMM Analysis:

CCMM's conditional scenario tree structure models successor threat branches: if a primary attack vector is mitigated, what is the probability that the adversary pivots to an alternative vector, and which vector is most likely given the adversary's documented capability profile?

Russian electronic warfare doctrine includes radiofrequency jamming and interference as standard tools against satellite communications links, with demonstrated capability across multiple operational theatres. Following the deployment of enhanced cyber mitigations by Viasat, the adversary retained full capability to conduct RF-based disruption against the same target class. The cyber mitigation deployment, by eliminating the preferred lower-cost, higher-deniability cyber approach, increases the probability of RF vector activation. This is a structural consequence of asymmetric mitigation: defending one attack surface does not reduce adversary motivation; it redirects adversary effort. **[AA]**

CCMM generates a probability of 63-68% for RF jamming and interference activation within 18 months of confirmed cyber mitigation deployment against the assessed target class. **[AA]**

Viasat's technical leadership confirmed in 2025 that this successor threat had materialised, consistent with the conditional prediction. **[CC: Via Satellite, Space Security Sentinel, 2025]**

Framework Comparison: No standard framework models conditional successor threats. ATT&CK catalogues RF interference as a technique when it is known and executed. STRIDE does not address the cyber-to-RF threat boundary. ISM has no mechanism for adversary vector pivot modelling.

Falsifiability Condition: If no confirmed RF interference or jamming activity against KA-SAT-class commercial satcom infrastructure is recorded within 18 months of confirmed cyber mitigation deployment, the successor threat probability revises downward to below 30%, triggering reassessment of adversary motivation and capability inputs.

6. The Low-Probability Domain

All five threat findings share a structural characteristic: at the time they were analytically derivable, standard frameworks would have assigned them negligible or zero analytical weight. None existed in any technique catalogue. None had been previously observed against this specific target class in this precise conditional configuration. In conventional threat assessment, this absence from the catalogue is treated as absence of the threat.

This is the defining feature of the low-probability, high-consequence domain. Events in this space are not random. They are conditionally structured, historically analogous, and detectable through cross-domain integration. They are not ordinarily visible to retrospective catalogue approaches precisely because they have not yet occurred at the time analysis is required, and their precondition sets span domains that no single framework integrates.

The 2022 KA-SAT attack did not emerge without precursor signals. It was the product of a Russian military doctrine applied consistently across two prior conflicts, executed through a governance gap created by a visible corporate transaction, against a dual-use network whose spillover characteristics were architecturally determined, by an actor whose identity was assessable from Day 1 through publicly available technical and geopolitical signals, with a successor threat profile that followed logically from the adversary's documented capability inventory.

The absence of prospective warning was not a failure of intelligence collection. It reflects the architectural scope of the frameworks that were applied. CCMM is specifically designed to address this gap, by operating in the domain where conditional structure, historical analogy, and cross-domain integration make prospective probability assignment possible, even when the specific event has not previously occurred.

7. Falsifiability Conditions: Summary

Table 7 summarises the observable indicators, revision triggers, and revised probability estimates applicable to each of the five threat findings. The presence of each indicator at the specified threshold constitutes a testable condition that, if confirmed, triggers downward revision of the associated probability assignment.

Table 7. Falsifiability conditions summary

Threat Finding	Falsifiability Indicator	Revision Trigger	Revised Probability
1. Conditional Timing Signal	No elevated GRU/Sandworm operational tempo within 30-day pre-escalation window	Confirmed absence of precursor indicators by independent assessment	Revises from 61-74% to below 35%
2. Governance Transition	Unified configuration audit completed within 90 days of transition commencement with independent verification	Audit completion confirmed	Governance uplift reduces by 60%
3. Civilian Spillover	Network segmentation score exceeds 0.75 on CCMM-Cyber criteria with independent routing verified	Independent segmentation assessment confirmed	Spillover probability reduces to below 25%
4. Attribution Ambiguity	AcidRain-variant code divergence from VPNFilter	Peer-reviewed technical malware	Attribution confidence revises to below 0.55 [RC]

Threat Finding	Falsifiability Indicator	Revision Trigger	Revised Probability
	lineage exceeds 40% on technical malware analysis	analysis confirms divergence threshold	
5. EW Succession	No confirmed RF interference or jamming against KA-SAT-class satcom within 18 months of cyber mitigation deployment	Absence confirmed by independent monitoring over the defined window	Successor probability revises to below 30%

The presence of these falsifiability conditions distinguishes CCMM assessments from analytical opinion. Each probability assignment is a testable claim: if the specified indicator is absent at the specified threshold, the assessment is revised downward. This mechanism transforms prospective threat assessment into an operationally accountable product. **[AA]**

8. Limitations, Validity, and Scope

8.1 Hindsight Bias

This paper presents a retrospective application of CCMM to a known historical case. The probability figures presented are derived from post-event evidence and are intended to illustrate the analytical output structure, not to claim that these exact figures would have been produced in a live pre-event application. In a prospective application, evidence availability and quality would differ, and probability ranges would be correspondingly wider. This is acknowledged explicitly and does not invalidate the methodological demonstration; it does, however, mean the paper's claims should be read as demonstrating analytical plausibility rather than proving prospective predictive accuracy.

8.2 Analyst Degrees of Freedom

The conditional branch structure and evidence-label assignments in this paper reflect the analytical judgements of a single analyst applying the CCMM framework. Different analysts may construct slightly different conditional structures or weight evidence labels differently. The public method layer in Appendix A is designed to constrain analyst degrees of freedom by specifying the ordinal logic governing branch activation and probability band generation. However, residual analyst discretion in input selection cannot be fully eliminated and should be acknowledged as a source of variance in any application.

8.3 Analogue Selection Sensitivity

The choice of Estonia 2007 and Georgia 2008 as historical analogues for Finding 1 is defensible against the available evidence, but other analysts may weight different analogues. The methodology requires that analogue selection be documented and justified, as it is in the calculation detail accompanying each finding. The sensitivity of the probability output to analogue selection is addressed in Section 8A.

8.4 Calibration Uncertainty

Probability assignments are calibrated through the proprietary CCMM scoring engine. The public method layer provides ordinal representations of the adjustment logic, but does not reproduce the precise calibration weights. This means external readers cannot verify the exact probability figures, only the direction and relative magnitude of adjustments. This is a known

limitation of proprietary methodology papers and is managed through the evidence labelling schema, which ensures every input claim is traceable to a source and quality class.

8.5 Omitted-Variable Risk

Cross-domain integration, while CCMM's primary advantage, also introduces the risk of analytical overreach: integrating geopolitical and technical signals across domains requires disciplined scope control to avoid importing analytical error from one domain into another. The five findings in this paper are each bounded by explicit conditional structures, evidence labels, and falsifiability conditions to mitigate this risk. Variables that could plausibly affect the probability outputs but were not available in the public evidence base, including classified intelligence assessments and internal Viasat security documentation, are not incorporated.

8.6 Single-Case Scope Limitations

The effective study scope is one case. KA-SAT is a strong demonstration case given the richness of its evidentiary record, but a single case is not sufficient to establish that CCMM reliably outperforms incumbent frameworks across the satcom threat environment more broadly. The paper does not make that claim. It makes the narrower claim that in this case, under this evidentiary boundary, CCMM produces prospective outputs that incumbent frameworks do not ordinarily generate. Multi-case validation, including positive replication cases, near-miss cases, and null cases where CCMM does not produce a differentiated output, is the necessary next step.

8A. Sensitivity Analysis

This section examines how the probability outputs for the five findings respond to changes in four analytical input conditions: altered analogue weights, downgraded evidence labels, reduced governance-transition weight, and removal of geopolitical intent signals. The purpose is to show that the outputs are directionally stable and analytically disciplined rather than impressionistic.

8A.1 Altered Analogue Weights (Finding 1)

The baseline probability range of 61-74% for Finding 1 is generated with documented HAC weights for the Estonia 2007 and Georgia 2008 analogues. The specific weight values are held as proprietary calibration parameters within the CCMM Technical Specification.

Under downweighted analogue scenarios, the probability range for Finding 1 reduces but remains above 50%, indicating moderate robustness to analogue selection variation. The direction of the output is stable; the magnitude is sensitive to analogue weighting. **[AA]**

8A.2 Downgraded Evidence Labels (Findings 1 and 4)

If the [CC]-labelled inputs for Finding 1 (GRU/Sandworm operational tempo signals, assessed from open-source reconnaissance reporting) are downgraded to [RC] to reflect a more conservative evidence assessment, the analytical confidence score reduces.

Evidence label downgrading produces material reductions in analytical confidence scores and probability ranges, but does not reverse the direction of any finding under the sensitivity conditions tested. The methodology's evidence-label architecture functions as designed: conservative labelling produces wider probability ranges and lower confidence classifications. **[AA]**

8A.3 Reduced Governance-Transition Weight (Finding 2)

The governance transition uplift applied in Finding 2 is the most architecture-specific parameter in the five findings.

Finding 2 shows the greatest sensitivity to parameter variation among the five findings. The adjusted attack-surface estimate is robust above the Gate A early warning threshold, but the Gate A designation is sensitive to weights near that threshold. This sensitivity should be disclosed in any commissioned engagement applying this finding to a specific operator context. **[AA]**

8A.4 Removal of Geopolitical Intent Signals (Findings 1 and 5)

Findings 1 and 5 both incorporate geopolitical intent signals as inputs. Removing these signals, to simulate a purely technical assessment without actor-intent inputs, tests the contribution of the geopolitical layer to the overall output.

The geopolitical intent layer contributes substantially to Findings 1 and 5. This is expected: CCMM's cross-domain advantage specifically lies in its ability to integrate actor intent signals with technical and organisational inputs. Removing the geopolitical layer reduces these findings to near-baseline estimates, confirming that the methodology's differentiated output is contingent on the quality of geopolitical intelligence inputs. This is an explicit design feature, not a limitation: prospective threat intelligence that incorporates actor intent is inherently dependent on the availability and quality of actor-intent evidence. **[AA]**

9. Conclusion

Standard frameworks remain valuable for classifying observed techniques, informing control selection, and supporting post-incident interpretation. The argument advanced in this paper is narrower and more specific: when used on their own, such frameworks are not ordinarily designed to generate prospective, probability-weighted, falsifiable threat assessments that integrate geopolitical posture, technical pathways, and organisational transition conditions within a single analytical structure.

Using the 2022 KA-SAT cyberattack as a retrospective demonstration case, this paper has shown how CCMM can be used to organise five categories of prospective threat finding: conditional timing signals preceding kinetic escalation, governance transition as attack surface, civilian infrastructure spillover in dual-use architectures, attribution ambiguity as an operational condition, and cyber-to-electronic-warfare succession following mitigation. In this case, those findings appear analytically accessible through a conditional cross-domain method even where incumbent frameworks, applied in isolation, produce primarily classificatory or retrospective outputs.

This case demonstrates the plausibility of conditional, falsifiable threat assessment in satellite communications. It does not, by itself, constitute definitive validation of framework performance across the wider satcom threat environment. The contribution of this paper is therefore methodological and demonstrative rather than conclusive.

For satellite communications operators, defence and intelligence organisations, and critical infrastructure providers operating in contested environments, the practical implication is that taxonomic and control-oriented frameworks may need to be complemented by an analytical layer capable of assessing whether protections are likely to hold under specific low-probability, high-consequence conditions. CCMM appears to offer a structured analytical complement in that role.

This paper should be read as the first CCMM cross-domain demonstration study in the satellite communications domain. The next step is broader multi-case validation, including comparative application to additional satellite communications and critical infrastructure incidents, sensitivity testing of calibration choices, and further refinement of the public method layer used to support external review.

10. References

Abeysekera, P. (2026). Conditional Consequence Mapping Methodology (CCMM): A Probabilistic Analytical Framework for Cross-Domain Threat Assessment. SSRN Abstract ID 6364078. Available at SSRN: https://ssrn.com/abstract=6364078

Abeysekera, P. (2026). Conditional Consequence Mapping Methodology (CCMM) [Data set]. Zenodo. https://doi.org/10.5281/zenodo.19382186

Betts, R.K. (1978). Analysis, War, and Decision: Why Intelligence Failures Are Inevitable. World Politics, 31(1), 61-89.

Boschetti, N., Gordon, N.G. and Falco, G. (2022). Space Cybersecurity Lessons Learned from the Viasat Cyberattack. ASCEND 2022. American Institute of Aeronautics and Astronautics.

Buchanan, B. (2020). The Hacker and the State: Cyber Attacks and the New Normal of Geopolitics. Cambridge, MA: Harvard University Press.

Council of the European Union. (2022, May 10). Russian cyber operations against Ukraine: Declaration by the High Representative on behalf of the European Union. Brussels: Council of the EU.

CyberPeace Institute. (2022). Case Study: Viasat Attack. Geneva: CyberPeace Institute. Available at: https://cyberconflicts.cyberpeaceinstitute.org

European Repository of Cyber Incidents (EuRepoC). (2023). Major Cyber Incident: KA-SAT 9A. Kerttunen, M., Schuck, K. and Hemmelskamp, J. 4 October 2023.

Falco, G. (2019). The Vacuum of Space Cyber Security. 2019 IEEE Aerospace Conference, Big Sky, Montana.

Falco, G., Viski, A. and Schmidt, S. (2022). Space Cybersecurity Lessons Learned from the Viasat Cyberattack. ASCEND 2022. American Institute of Aeronautics and Astronautics.

Grabo, C.M. (2004). Anticipating Surprise: Analysis for Strategic Warning. Washington DC: University Press of America.

Guerrero-Saade, J.A. and van Amerongen, M. (2022). AcidRain: A Modem Wiper Rains Down on Europe. SentinelOne Threat Intelligence. 31 March 2022.

Heuer, R.J. and Pherson, R.H. (2014). Structured Analytic Techniques for Intelligence Analysis. 2nd ed. Washington DC: CQ Press.

Kostyuk, N. and Zhukov, Y.M. (2019). Invisible Digital Front: Can Cyber Attacks Shape Battlefield Events? Journal of Conflict Resolution, 63(2), 317-347.

Libicki, M.C. (2009). Cyberdeterrence and Cyberwar. Santa Monica, CA: RAND Corporation.

Mandel, D.R. and Barnes, A. (2014). Accuracy of Forecasts in Strategic Intelligence. Proceedings of the National Academy of Sciences, 111(30), 10984-10989.

MITRE Corporation. (2023). ATT&CK for Space. MITRE ATT&CK Knowledge Base. Available at: https://attack.mitre.org

Rid, T. and Buchanan, B. (2015). Attributing Cyber Attacks. Journal of Strategic Studies, 38(1-2), 4-37.

Rinaldi, S.M., Peerenboom, J.P. and Kelly, T.K. (2001). Identifying, Understanding, and Analyzing Critical Infrastructure Interdependencies. IEEE Control Systems Magazine, 21(6), 11-25.

Tetlock, P.E. and Gardner, D. (2015). Superforecasting: The Art and Science of Prediction. New York: Crown Publishers.

UK Government. (2022, May 10). Russia behind cyber-attack with Europe-wide impact an hour before Ukraine invasion. London: UK Government.

US Department of State. (2022, May 10). Attribution of Russia's Malicious Cyber Activity Against Ukraine. Statement of Secretary Antony J. Blinken. Washington DC: US Department of State.

Via Satellite / Space Security Sentinel. (2025). Three Years Post KA-SAT Attack, Viasat Exec Talks Lessons Learned on Cybersecurity Posture. Interview with Phil Mar, VP/CTO Engineering, Viasat Government.

Viasat Inc. (2022, March 30). KA-SAT Network Cyber Attack Overview. Carlsbad, CA: Viasat Inc. Available at: https://viasat.com

Appendix A: Public Method Layer - Ordinal Pseudocode

This appendix provides a non-proprietary representation of the CCMM analytical method used to generate the probability outputs in this paper. It is provided in ordinal pseudocode form to support external review of the analytical structure without disclosing the proprietary calibration weights and threshold values held within the CCMM Technical Specification. Precise calibration values remain proprietary to GABEY Consulting Pty Ltd (ACN 121 511 055).

A.1 Branch Definition

Step 1: Define the scenario domain and target class.

- Identify the threat category being assessed and its relevant domain (geopolitical, technical, or organisational).
- Establish the conditional structure: what preconditions must be met for the threat to activate.
- Assign an activation threshold: the minimum set of conditions required before a probability output is generated.
- If activation conditions are not met, do not generate a probability output. Flag the finding for evidence development.

A.2 Evidence Ingestion and Labelling

Step 2: Ingest evidence and apply evidence labels.

- FOR EACH claim in the evidence set:
- IF source is official government statement or formally verified finding: label [OF].
- IF source is corroborated by multiple independent sources without official verification: label [CC].
- IF source is analyst inference from available evidence without independent confirmation: label [AA].
- IF source is a single unverified report: label [RC].
- Apply the label at the point of initial drafting. Do not assign a higher label than the source quality warrants.

A.3 Precondition Convergence Assessment

Step 3: Assess precondition convergence.

- FOR EACH precondition in the conditional branch:
- IF precondition is supported by [OF] or [CC] evidence: assign HIGH ordinal weight.
- IF precondition is supported by [AA] evidence: assign MODERATE ordinal weight.
- IF precondition is supported by [RC] evidence: assign LOW ordinal weight.
- Convergence score = ordinal sum across all required preconditions.
- IF convergence score is HIGH: proceed to probability band generation.
- IF convergence score is MODERATE: note analytical uncertainty in output and widen the probability range.

- IF convergence score is LOW: do not generate probability output; flag for evidence development.

A.4 Historical Analogue Comparison (HAC)

Step 4: Apply historical analogue comparison.

- IDENTIFY analogues from the historical record where: actor class is comparable; target class is comparable; operational context is comparable.
- WEIGHT analogues by: recency (more recent analogues receive higher weight); specificity (more specific analogues receive higher weight); outcome verifiability (analogues with confirmed outcomes receive higher weight).
- APPLY analogue weight as an ordinal adjustment to the convergence score.
- DOCUMENT the analogue selection rationale. Analogue selection must be explicit and justified.

A.5 Probability Band Generation

Step 5: Generate probability band.

- IF convergence score plus analogue adjustment is HIGH: assign upper probability band (broadly in the 60-80% range).
- IF convergence score plus analogue adjustment is MODERATE-HIGH: assign mid-upper band (broadly in the 40-65% range).
- IF convergence score plus analogue adjustment is MODERATE: assign mid band (broadly in the 25-45% range).
- IF convergence score plus analogue adjustment is LOW-MODERATE: assign lower band (broadly in the 10-30% range).

Note: Exact band boundaries are set by the proprietary calibration layer within the CCMM scoring engine and are not disclosed in this public method representation. The ordinal bands above are approximate ranges for interpretive reference only.

A.6 Revision Trigger Specification

Step 6: State the revision trigger.

- FOR EACH probability output, DEFINE the observable indicator whose absence by a threshold date reduces the estimate.
- SPECIFY the direction and approximate magnitude of revision.
- ENSURE the revision trigger is independent of the initial estimate and not circular.
- ENSURE the revision trigger is observable by an analyst without access to classified information, unless the assessment is produced in a classified context.

A.7 Evidence-Labelled Reporting

Step 7: Report with evidence labelling.

- Label the final output with the dominant evidence class.
- IF primarily [OF]-anchored: state the [OF] basis explicitly.
- IF primarily analytical inference: label [AA] and note the [OF] or [CC] anchoring claims that provide factual grounding.

- IF mixed: label [AA] and itemise the evidence class of each input claim.
- DO NOT label [AA] outputs as confirmed. Preserve analytical distance between assessment and verified fact.
- Include all source citations against individual claims at the point of assertion, not only in the reference list.

This public method layer is provided for external review of the analytical architecture used in this demonstration study. It does not constitute a disclosure of proprietary calibration weights, threshold values, or scoring engine parameters. Those elements are held as proprietary intellectual property of GABEY Consulting Pty Ltd (ACN 121 511 055) and are available to licensed practitioners under the CCMM Technical Specification.