[security-wg] Policy Proposal 2026-02 - Response to Excedo Networks

Thank you for your exceptionally thorough and operationally grounded analysis. Input of this caliber from an active Security Operations Center and Digital Crimes Unit is invaluable.

Part 1: Scope & Registry Integrity (Points 2 & 14)

Because your response covers a wide range of legal and technical layers, I would like to address your points systematically, starting with the core contractual and registry integrity questions raised in Point 2 and Point 14:

1. Registry Integrity & Struck-Off Entities (Your Point 2)

Your observation regarding active ASNs operating under legally dissolved or struck-off shell companies highlights a severe operational frustration for threat investigators. You are entirely correct that a public non-compliance flag is an insufficient tool if the registered counterparty no longer exists as a legal entity.

However, addressing the legal existence of a member falls strictly outside the scope of this specific abuse mitigation policy. The RIPE community already possesses a dedicated, authoritative framework for this explicit scenario: RIPE-858 (Closure of Members, Deregistration of Internet Resources and Legacy Internet Resources).

Under RIPE-858, a resource holder operating under a dissolved entity or maintaining fraudulent contact data is already subject to immediate contract closure and permanent resource deregistration on the grounds of providing "Untruthful information." If an entity has ceased to exist, the RIPE NCC should not be applying the temporary, reversible restrictions proposed in 2026-02; they should be executing standard RIPE-858 contract termination.

Even though it is outside the scope of this policy, it is a related matter that the RIPE NCC must address. I would suggest where there is credible evidence that a legal entity who is registered as a RIPE NCC member and holds internet resources issued by the RIPE NCC to that legal entity, no longer exists, the RIPE NCC needs to establish a procedure, if one doesn't already exist, to report this matter and provide the credible evidence to them. Please open this matter on a separate thread if you want to follow it up.

2. Contractual Arrangements & Administrative Burden (Your Point 14)

You suggest that the policy needs to explicitly distinguish between direct members, sponsored resource holders, and other contractual variations.

From a policy design perspective, this framework intentionally treats all RIPE NCC-issued address space uniformly. The underlying contractual variance does not alter the core obligation: address space distributed by the registry must not systematically facilitate the infrastructure cybercrime defined in Article 1.

The primary administrative burden we must avoid is forcing RIPE NCC staff to perform subjective, judicial evaluations of network traffic. By relying strictly on the binary trigger of an Approved Infrastructure Blocklist combined with manual reporting verification, the RIPE NCC

operates purely as an administrative registry auditor. They verify the data presence, relay the notice to the official contract address on record, and programmatically apply the temporary, completely reversible registry restrictions until the malicious activity is confirmed to have stopped.

Part 2: Enforceable Obligations, Threat Definitions, and Subscriber Protections (Points 1, 3, & 4)

Continuing our technical alignment, I would like to address your specific semantic and definitional critiques raised across Point 1, Point 3, and Point 4:

3. The Dual Nature of Enforceable Obligations (Your Point 1)

You argue that a distinction must be made between a resource holder's response to a report and the underlying presence of abuse as stated in Article 2.1.

In reality, these two obligations are fundamentally intertwined. This policy framework establishes that resource holders are contractually accountable for systematically failing to address abuse, because under Article 8 (Liability) of the RIPE NCC Standard Service Agreement (SSA), members explicitly agree that: "The Member shall also be liable for all aspects of its use and all that ensues from its use of the Internet Number Resources."

This contractual liability undeniably covers how and to whom an LIR assigns its address space, and how those End Users utilise it. The member is legally responsible for any regular abuse originating from resources under their custody. How they choose to monitor, investigate, and document their internal remediation is entirely their operational responsibility. The presence of regular, non-responsive infrastructure abuse is the definitive proof of a failure to meet that obligation.

4. Technical Thresholds and Boundary Definitions (Your Point 3)

You noted that terms like "sustained," "systemic," and "dedicated" require clearer definitions to prevent platform gaming. I accept your critique regarding textual tightening. To achieve complete clarity, Article 1.1 will be adjusted in the next iteration to read:

"1.1 For the purposes of RIPE Policy, 'Network Abuse' is defined as regular, infrastructure-level malicious operations, specifically including..."

To satisfy your request for explicit operational definitions under this framework:
- "Regular malicious operations" means any recurring, planned, or systematic harmful activity executed by malware, threat actors, or automated scripts against a computer system, network, or user.
- "Systemic" means the malicious activity is pervasive, deeply embedded, and institutional, rather than an isolated collection of one-off, compromised end-user nodes.
- "Dedicated" (in the context of DDoS command-and-control) means the underlying servers or software components are exclusively set up and utilised for managing malicious infrastructure, rather than sharing general-purpose tasks.

- "Systematically facilitate" means to regularly, efficiently, and structurally enable these threat classifications.

Regarding attribution, you state the policy must require evidence connecting the responsible relationship to "identify the party controlling the abusive service."

This misinterprets the charter of an Internet registry. As the community routinely reminds us, the RIPE NCC is not the internet police. It is not the RIPE NCC's job to investigate, gather evidence, or catch criminals—that is the job of state law enforcement. This policy does not require the RIPE NCC to carry out investigations. The registry's only job is to audit the asset's data state. The burden of investigation sits entirely on the resource holder; they are the ones who possess the direct customer relationship and know exactly what service is being provisioned on their blocks.

Finally, the threat classifications listed under Article 1.1 are a starting suggestion. The final content of this list belongs entirely to the community, and I welcome Working Group consensus to expand or refine these items.

5. Residential and Dynamic Subscriber Pool Scope (Your Point 4)

In Article 1.3, the framework intentionally excludes consumer broadband and dynamic subscriber pools to protect casual end-user compromises from contractual penalties. You noted that the reference to the Spamhaus PBL requires refinement, as it can include static infrastructure.

I welcome your feedback on this boundary. The RIPE community is full of world-class routing and filtering experts. If the Working Group can offer a tighter, more robust technical definition for verifying and classifying residential access infrastructure, I will be grateful for that assistance and will gladly incorporate it into the next version 2.0.

Part 3: Reporting Eligibility, Evidentiary Parameters, and Trigger Telemetry (Points 5, 6, & 7)

Moving on to your structural concerns regarding reporting channels and data validation workflows in Point 6, Point 5, and Point 7:

6. Reporting Eligibility & The RIPE Database-Native CSIRT Gate (Your Point 6)

You rightly point out that restricting reporting eligibility strictly to affected network operators or National CERTs under Article 1.4 creates an operational bottleneck. It inadvertently excludes premier, frontline incident-response organisations and private Digital Crimes Units who possess high-signal forensic data but are not the direct victims of the specific routing announcement.

To resolve this without opening the RIPE NCC Support queue to unverified third-party ticket spam, we can utilise a long-standing, native RIPE Database architecture: the irt (Incident Response Team) object.

The irt object was engineered specifically to define a verified Computer Security Incident Response Team (CSIRT) handling serious infrastructure-level threats. For an independent security firm, threat intelligence network, or research organisation to qualify as an authorised reporter under Article 1.4, the framework will require them to maintain an active, fully validated irt object in the RIPE Database. Because irt object creation requires formal due diligence or global accreditation via trust-brokers like the Trusted Introducer (TI) service, this establishes an immediate, accountable baseline for reporting entities.

7. Evidentiary Parameters & Registry Neutrality (Your Point 5 and Point 7)

You lay out a detailed 7-point baseline for an "actionable evidence package," including collection methodologies and persistence chronologies, suggesting that the policy should distinguish between reporter authentication and evidentiary sufficiency. Furthermore, in Point 7, you query how partial days, intermittent observations, and combined lists are counted.

This requires an immediate clarification of scope to protect the core neutrality of the RIPE registry. If this framework requires the RIPE NCC to collect, review, and evaluate complex forensic network logs, the RIPE NCC would be forced to make semi-judicial decisions on the technical validity of data trails. This completely contradicts the administrative design of the framework.

To create a completely workable policy that avoids placing the RIPE NCC in a judicial review role, we must rely on fixed, binary RIPE registry criteria. I have chosen the criteria to be a minimum of an active listing for a cumulative 14 days across one or more Approved Infrastructure Blocklists within a 30-day rolling window. To guarantee freshness, the authenticated notification must be submitted to the RIPE NCC within 30 days of the last recorded listing day.

As you rightly point out, a cumulative 14-day listing may not perfectly equate to 14 days of active, continuous malicious traffic, but it is the uniform administrative metric we are setting. How individual observations or partial days are calculated within those metrics is strictly a technical matter determined by the telemetry of the Approved Blocklist operator, not the RIPE NCC.

The RIPE NCC does not undertake investigations or evaluate evidence. Under the signed SSA, it is the resource holder's absolute responsibility to ensure their address space does not engage in abusive behaviour. It is for the LIR to investigate and consider any forensic reports sent to them. As long as the prefixes remain listed on an Approved Blocklist, this framework considers the resource holder to have contractually failed to resolve the matter, regardless of whether they claim they failed to receive individual end-user reports.

If an accredited security firm possessing an active irt object wishes to compile deeper forensic packages, that evidence should be delivered directly to the resource holder (the LIR) during the 7-business-day grace period to facilitate their internal investigation under their Article 8 SSA obligations. The RIPE NCC's role is kept strictly to a binary administrative audit of registry records and blocklist presence, preventing the RIPE registry from being pulled into a judicial review role.

You raise a critical technical challenge regarding evasion: "activity may continue through rotating addresses without any individual address reaching the threshold."

To close this loophole without increasing the RIPE NCC's administrative burden, the policy tracks telemetry at the registered RIPE Database block layer. That may be an assignment, aggregation, or allocation containing the IP address(es). When a malicious operator utilises agile IP rotation across a block of addresses, Approved Infrastructure Blocklists (such as the Spamhaus DROP list) automatically aggregate the rotating threat pattern and list the encompassing block. If the component addresses or the covering block collectively trigger the 14-day threshold within the 30-day window, the RIPE NCC executes the standard binary check on that block. This ensures that fast-fluxing or IP-hopping cannot be utilised to evade contractual registry accountability.

Part 4: Route Origin Authorisation (ROA) Technical Alignment & Compliance Pathways (Point 12)

Next, I would like to address your technical critique regarding Route Origin Authorisations (ROAs) and compliance restoration pathways raised in Point 12:

8. ROA Removal and Alternative Compliance Review Windows (Your Point 12)

Your technical analysis of Article 3.6 regarding Route Origin Authorisation (ROA) deletion is exceptionally sharp. As you correctly point out, deleting a ROA does not drop a BGP routing announcement or take an abusive service offline; depending on remaining covering aggregate objects, it merely shifts the validation state to "Not Found" or "Invalid." This does not independently prove that the infrastructure abuse has stopped, and it could create an unintended incentive for non-compliant operators to weaken their routing-security protections simply to satisfy a registry restoration loop.

Therefore, I accept your recommendation to remove ROA deletion as a standalone proof of compliance.

However, you raise an equally vital operational point: if we rely exclusively on automated delisting, we create a dangerous deadlock for the LIR if a listing becomes stale, or if a delisting process is delayed by a protracted dispute with an external blocklist operator. An effective framework must recognise verified mitigation even when an external listing remains temporarily stagnant.

To prevent this deadlock without imposing an unworkable technical burden, this framework will not expect the RIPE NCC software to continuously and programmatically scan external blocklists for removals. Compliance review will be manually instigated by the resource holder responding directly to the open RIPE NCC compliance ticket.

Version 2.0 will introduce a dual-pathway restoration workflow to handle stale or disputed listings:
- Pathway A (Standard): The LIR requests a review by demonstrating that the prefix has successfully dropped off the Approved Infrastructure Blocklist.

- Pathway B (Alternative): If a listing remains stale or disputed, the LIR may submit a ticket update documenting that verified mitigation has occurred at the infrastructure layer (such as demonstrating that the malicious downstream interface has been completely terminated or the specific aggregate route pulled).

The RIPE NCC staff will then execute a binary administrative check, within one working day of the ticket update, to confirm the infrastructure isolation is complete, and then restore registry agility without pulling the RIPE registry into a judicial evaluation of the blocklist dispute itself.

Part 5: Delivery Procedures and Contractual Handoffs (Point 9)

Next, I would like to address your operational concerns regarding delivery tracking, contact targets, and grace period execution windows raised in Point 9:

9. Contractual Contact Responsibility and Operational Delivery Hooks (Your Point 9)

You raise an important operational point regarding the targets of formal notifications. While ordinary abuse handling may occur at the local network layer, once a failure triggers a formal compliance notice, the escalation must engage both the technical operational teams and the authoritative legal entities responsible for the contract.

I accept your recommendation that technical and abuse contacts must not be bypassed. To achieve this workflow balance without weakening the contractual timeline, Article 3.3 and Article 3.4 will be rewritten in Version 2.0 to establish an airtight delivery protocol:

"3.3 The administrative delivery protocol requires that the RIPE NCC shall send a formal remediation notice by means of email to the authoritative legal, administrative, and billing contacts the resource holder is contractually required to maintain under their signed RIPE NCC Standard Service Agreement (SSA). If the abuse involves multiple IP addresses distributed across multiple registries under legally affiliated accounts or consolidated billing structures, copies of the remediation notice will be forwarded to all associated accounts. Copies will also be sent to the abuse-c and tech-c contacts of the appropriate allocation object related to the abusive addresses.

3.4 A resource holder is granted a strict grace period of 7 business days from the date the remediation notice is sent to achieve compliance. Compliance is achieved when all identified abusive prefixes successfully drop off the Approved Infrastructure Blocklists, or when the resource holder takes verified administrative action to utilise the alternative pathway to isolate the abusive prefixes. The resource holder is contractually required to maintain accurate and correct contact details, for official notifications and within their RIPE Database objects, allowing the RIPE NCC to be able to contact them. So once the remediation notice is sent it can be assumed that it will be delivered."

Regarding your suggestion for defined retries, alternative delivery methodologies, and fallback grace period triggers: these mechanisms are explicitly omitted from the framework.

Under the signed SSA and current RIPE Database business rules, it is the sole, non-delegable contractual responsibility of the resource holder to ensure their registered email pathways are fully functional, monitored, and available to receive official notifications. The RIPE registry cannot accommodate compliance delays based on an operator's internal mail infrastructure failures or unmonitored queues; the strict 7-business-day clock triggers the moment the transmission is executed to the contacts on record.

10. Operational Timeline Standardisation and Anonymised Appeal Routing (Your Point 10 and Point 13)

You note in Point 10 that the initial draft contains a textual inconsistency between "business days" and "working day," and you argue in Point 13 that the framework lacks an explicit pathway for a resource holder to challenge evidence, file an independent appeal, or secure a procedural pause during an active review.

To achieve a clean, standardized timeline across the RIPE service region, Version 2.0 will completely standardise on the term 'business days' rather than 'working days.' In several jurisdictions across our region, working days can routinely include Saturdays, but for the RIPE NCC, the clock will be bound strictly to its official corporate calendar.

To resolve your concerns regarding procedural safeguards without pulling the RIPE registry into a non-neutral, judicial role, we must implement a highly structured, automated compliance workflow. The RIPE NCC must not be required to evaluate technical responses, determine what constitutes "verifiable containment," parse generic assurances, or act as an appellate court judging forensic case files.

Instead, all substantive evaluations of evidence, containment timelines, and ongoing mitigation adjustments will be addressed directly to the notifier. The notifier is the entity possessing the live threat telemetry and is the only party capable of assessing an operator's counter-arguments or technical remediation steps. To protect independent investigators from corporate intimidation or retaliation, the framework must guarantee total anonymity for the reporting party.

Therefore, Version 2.0 will introduce an automated, anonymised Subject Line Mirror protocol within the RIPE NCC ticketing queue:

- The "APPEAL" Trigger: If a resource holder wishes to question the blocklist evidence or document their containment actions, they must reply to the open RIPE NCC ticket with the word 'APPEAL' in the subject line. The RIPE NCC ticketing engine will automatically parse this string and mirror the text directly to the notifier's registered irt mailbox.
- The Anonymised Mirror Return: A response from the notifier with the word 'APPEAL' remaining in the subject line will be programmatically routed back to the resource holder with all identifying contact details of the notifier completely scrubbed.
- The "CANCEL" Trigger: If the notifier determines that the investigator's evidence is invalid or that sufficient, appropriate mitigation has occurred at the infrastructure layer, they may reply to the ticket with the word 'CANCEL' in the subject line. The RIPE NCC software will instantly lift any applied administrative restrictions, purge any non-compliant status attributes, and close the compliance ticket—even if the underlying addresses have not yet formally dropped off the external Approved Blocklist.

- The "DELAY" Trigger: If the notifier considers that the resource holder is actively cooperating but requires an operational window to complete deep remediation, they may reply with the word 'DELAY' in the subject line. If no intermediate restrictions have yet been applied, the RIPE NCC ticketing software will automatically append an additional 7 business days to the original grace period.

While this protocol utilizes multi-layered message routing, it is explicitly designed to allow the RIPE NCC to automate the entire compliance pipeline at the software layer while maintaining absolute operational privacy for threat intelligence contributors.

Finally, regarding your comment on "urgent operational escalation" for cases of serious, continuing harm: this falls completely outside the scope of an administrative registry policy. The RIPE NCC are not the internet police, and it is not the registry's job to mitigate live kinetic emergencies. If a threat vector represents an immediate public safety crisis, the matter must be referred directly to actual state law enforcement agencies.

Part 6: Support Document Architecture and Outcome Metrics (Points 11, 8, & 15)

To conclude our line-by-line review, I would like to address your final points regarding restriction transparency, blocklist governance, and outcome metrics raised in Point 11, Point 8, and Point 15:

11. Concrete Operational Restrictions (Your Point 11)

You note that Article 3.5 cannot be fully evaluated in a vacuum without reviewing the exact technical parameters of the restrictions themselves. I agree completely.

The granular intermediate registry restrictions you highlighted are fully detailed within the separate RIPE Network Abuse Policy Support Document, which I launched on the mailing list for parallel discussion alongside the main policy articles. The complete text can be viewed directly in the public archive here: https://mailman.ripe.net/archives/list/security-wg@ripe.net/message/ZAHSW7B6VYCJCN46JCAG3SUPGKNMEGE3/

Please note that because standard list indexing can alter clean spacing, much of the visual layout has been flattened in the raw body text. If you click on the 'attachment' button at the bottom of that archive message, you can download the complete standalone document in HTML format with its proper engineering layout and formatting intact.

12. Governance of Approved Blocklists (Your Point 8)

You rightly state that the criteria and approval of blocklists constitute a substantive part of this policy's long-term integrity. The baseline registry of Approved Blocklists is also contained natively within Section 2 of the parallel Support Document.

All technical details, data entry standards, and telemetry criteria regarding the inclusion or removal of these feeds will be evaluated and decided exclusively by consensus within the

Security Working Group. Because the Support Document is a community-governed RIPE document rather than rigid text-law, the Working Group maintains the absolute authority to revise, audit, or withdraw approval from any dataset at any time as threat landscapes shift.

13. Measuring Effectiveness Through Outcomes (Your Point 15)

Your suggestions regarding tracking specific outcome metrics—such as the time to a substantive response, verified containment intervals, recurrence rates, and collateral impact logs—are excellent and highly interesting. Because these operational metrics relate directly to internal registry data tracking and support workflows, I will leave these suggestions entirely for the RIPE NCC staff to consider within their formal Implementation Assessment.

Summary of Objectives

Michael, with the extremely valuable points you have raised within this critique and my detailed responses across these sections, I believe that together we have successfully addressed your entire list of foundational objectives.

I look forward to incorporating these technical boundaries into a refined Version 2.0 draft, and I welcome the continued collaboration of yourself and others as we deliver a robust, secure, and operationally defensive framework to the registry.