A National-level Real-time DDoS Mitigation Architecture Coordinated by DNSC at Romanian Internet Exchange Points

Alexandru TRANDAFIRESCU, Nicolae ȚĂPUȘ
Faculty of Automatic Control and Computer Science, National University of Science and Technology POLITEHNICA Bucharest (UNSTPB), Romania
alex@qsign.ro, nicolae.tapus@upb.ro

Abstract: This paper provides a critical analysis of Romania's Cyber Security Strategy (2022–2027) in the context of the rapidly evolving hyper-volumetric DDoS threat landscape, amplified by offensive Artificial Intelligence capabilities. By examining global and national attack statistics from authoritative sources (Cloudflare, NETSCOUT, DNSC) covering 2024–2025, the paper identifies a significant gap between existing legislative and procedural frameworks and operational reality. A real-time DDoS mitigation mechanism is proposed, leveraging automated detection at the edge router level combined with coordinated upstream filtering at Internet Exchange Points (IXPs) and international gateways, orchestrated by DNSC as the national authority. The proposed architecture comprises a REST API, end-user pre-registration with IP-ownership verification and out-of-band backup IP provisioning, and ISP/gateway pre-registration as filtering endpoints. A comparative analysis with commercial scrubbing services (starting from 700 EUR/month) shows that the proposed model eliminates architectural complexity, privacy exposure and vendor lock-in while achieving superior national scalability at near-zero marginal cost.
Keywords: DDoS mitigation, Internet Exchange Point, DNSC, NIS2, cyber security strategy, real-time incident response, upstream filtering, IP ownership verification, REST API.

INTRODUCTION

The current normative framework, defined by the publication in the Official Gazette no. 2 bis of 3 January 2022, of Romania's Cyber Security Strategy for the period 2022–2027 (Government of Romania, 2022), represents a fundamental pillar in aligning the national security posture to European standards. The strategic document identifies the cyber risks and threats anticipated at the time of drafting, offering a roadmap for national resilience.

https://doi.org/10.54851/v8i2y202603

Subsequently, Directive (EU) 2022/2555 – NIS2 (European Parliament & Council, 2022) – was transposed into Romanian legislation, extending the scope of application and tightening the security requirements for operators of essential services.

From the perspective of 2026, however, a discrepancy emerges between the proactive measures established in 2022 and current operational reality. The major disruptive factor is the proliferation and democratisation of Artificial Intelligence (AI) tools for offensive purposes (ENISA, 2024). This evolution has catalysed an unprecedented escalation of cyber attacks, indiscriminately affecting data infrastructure ranging from small data centres to major critical infrastructure.

The contribution of this paper is twofold. First, it provides a comparative analysis of global and national DDoS attack data to substantiate the urgency of updating technical and procedural measures. Second, it proposes a practical, cost-effective, and technically feasible mechanism for real-time DDoS mitigation at national level, based on automated detection, DNSC-validated coordination, and upstream filtering at IXP/gateway level using standard nftables rules deployable on any Linux-based router. The system is intentionally agnostic to the operating systems and hardware platforms used by participants, in order to maximise its deployability across the existing national infrastructure.

COMPARATIVE ANALYSIS OF DDOS ATTACKS: GLOBAL VS. NATIONAL CONTEXT

Global threat panorama

Recent statistics highlight a significant growth in hyper-volumetric attacks. In October 2024, Cloudflare mitigated an attack reaching 5.6 Terabits per second (Tbps), surpassing previous records of 3.8 Tbps and 4.2 Tbps recorded earlier in the same year (Cloudflare, 2025); the incident was generated by a Mirai-variant botnet exploiting IoT device vulnerabilities. The year 2024 marked a 53% increase in attacks blocked by Cloudflare, totalling 21.3 million incidents (Cloudflare, 2025), while NETSCOUT reported more than 8 million incidents in the first half of 2025 alone (NETSCOUT, 2025). A new record of 3.15 billion packets per second (pps) was reported by Global Secure Layer in August 2024 (GSL, 2024); this metric is critical because it targets the exhaustion of network equipment CPU processing resources, not merely bandwidth saturation. A predominance of “hit-and-run” attacks is observed, where approximately 90% of incidents last under 10 minutes (Cloudflare, 2025); their tactical purpose ranges from defence testing and harassment to their use as diversions for more complex intrusions such as ransomware deployment (Microsoft, 2024).

Impact on Romanian infrastructure

At national level, aggregated data from the National Directorate for Cyber Security (DNSC) and operators such as Orange confirm Romania’s alignment with the global trend, with locally concerning specificities. In 2024, DNSC and Orange Business monitored over 27 million security events (DNSC, 2025). The largest recorded attack reached 184.39 Gbps (NETSCOUT, 2025). While this value is inferior to global records expressed in Tbps, an attack of 184 Gbps is sufficient to render non-functional the majority of public institutions and companies in Romania that do not benefit from enterprise-level traffic scrubbing solutions. Peak flows of 66.1 million packets per second (Mpps) were also recorded (NETSCOUT, 2025); such volume is capable of blocking conventional firewalls and routers through hardware resource exhaustion, even in the absence of complete bandwidth saturation. Similar to the global trend, attacks

in Romania are short (under 30 minutes) but frequent, with the primary objective of service unavailability and reputational damage rather than permanent infrastructure destruction.

EXISTING COMMERCIAL DDOS MITIGATION SOLUTIONS ON THE ROMANIAN MARKET

Commercial DDoS mitigation services already exist on the Romanian market. These solutions, while effective, come at significant cost and are primarily accessible to large enterprises and critical-infrastructure operators. DIGI Romania (RCS & RDS), the country's largest telecommunications provider, offers enterprise-grade DDoS traffic scrubbing packages that clean traffic before it reaches the customer's equipment, starting from approximately 700 EUR per month per protected IP prefix (DIGI, 2025). Similar services are offered internationally by providers such as Cloudflare (Magic Transit), Akamai (Prolexic), and Radware, with prices scaling from several hundred to tens of thousands of euros monthly depending on bandwidth and attack-capacity guarantees (Radware, 2024). At national level, providers operating within Romanian data centres (Voxility, M247, NXDATA-1) offer integrated DDoS protection through on-path or out-of-path scrubbing devices, typically based on A10 Networks or Arbor Networks equipment (DedicatServer, 2024). While technically competent, these solutions protect only the specific hosted infrastructure and do not offer a systemic, national-level mitigation capability.

Beyond industry and governmental reports, the recent academic literature confirms the same trajectory. Apostu et al. (2025) survey artificial-intelligence-based techniques for detecting and mitigating DDoS attacks, arguing that automated, machine-assisted response is becoming a prerequisite for defending against machine-generated attack traffic. In the same direction, Vevera et al. (2025), in a recent study published in the Romanian Cyber Security Journal, analyse the threat landscape of domain name infrastructures and identify DNS amplification and NXDOMAIN-flood attacks as persistent vectors that intensify volumetric DDoS assaults, concluding that no single point solution suffices and that layered, coordinated defence mechanisms are required. These findings reinforce the case for the infrastructure-level, nationally coordinated filtering approach proposed in this paper.

The critical observation is that these commercial solutions address the symptom rather than the root cause. Each organisation must independently procure and manage its own DDoS protection, creating an uneven security landscape where well-funded entities are protected while smaller organisations, public institutions, healthcare providers, and educational institutions remain vulnerable.

Architectural complexity of current scrubbing services

Cloud-based scrubbing services such as Cloudflare Magic Transit, Akamai Prolexic, or Radware DefensePro require the protected entity to fundamentally alter its network architecture. Under the DNS-based redirection model – the most common deployment – the customer must delegate its authoritative DNS to the scrubbing provider, meaning that all DNS resolution for protected domains is handled by third-party infrastructure (Cloudflare, 2025). For BGP-based scrubbing, the provider announces the victim's IP prefix via BGP, attracting all inbound traffic to the scrubbing centre, cleaning it, and forwarding only legitimate traffic back to the origin via GRE tunnels or direct interconnects (Radware, 2024). This introduces multiple additional points of failure: GRE tunnel instability, increased latency through the scrubbing path, and dependency on the availability and correct configuration of the scrubbing provider's

infrastructure. In practice, the operator must also maintain compatibility with the provider's SSL/TLS termination requirements (which may involve sharing private keys or using provider-issued certificates) and manage split-horizon DNS for services that need to bypass the scrubbing layer.

Privacy and data sovereignty concerns

A frequently overlooked consequence of routing traffic through third-party scrubbing centres is the privacy implication: all network traffic – including legitimate user data, API calls, authentication tokens, and potentially sensitive communications – transits through infrastructure controlled by a foreign commercial entity. For Romanian public institutions, healthcare providers, and entities handling personal data under GDPR (European Parliament & Council, 2016), this creates a problematic data flow where citizen data passes through servers outside national jurisdiction, even if only transiently during the scrubbing process. Furthermore, certain providers (notably Cloudflare) require operators to use their DNS resolvers and proxy infrastructure, effectively positioning the scrubbing provider as a man-in-the-middle for all HTTPS traffic. While robust security controls are in place, the fundamental architectural reality is that decrypted traffic is inspected at a third-party facility – a requirement that may conflict with data-minimisation principles under Article 5(1)(c) of GDPR and with sectoral regulations for healthcare, financial, and governmental data processing (ENISA, 2024).

Scalability and national infrastructure load

From a national infrastructure perspective, current scrubbing solutions are inherently non-scalable in aggregate. Each micro data centre that adopts a cloud scrubbing service must independently provision bandwidth to forward its traffic to the scrubbing destination and receive cleaned traffic back. During a coordinated DDoS campaign targeting multiple Romanian entities simultaneously, each victim's traffic must travel to a scrubbing centre often located outside Romania (Frankfurt, Amsterdam, London), be processed, and return – effectively doubling or tripling the volume of traffic transiting Romania's international gateways. The mitigation mechanism itself thus increases the load on the very infrastructure it aims to protect. Moreover, a micro data centre with a typical 10 Gbps link must provision sufficient bandwidth not only for legitimate traffic but also for the overhead of the scrubbing tunnel; under a 5 Gbps DDoS attack such a link has effectively zero capacity remaining for legitimate traffic, because the attack traffic must still traverse the last-mile link before being filtered.

Comparative advantage of the proposed model

In contrast, the DNSC-coordinated whitelist/blacklist model proposed in the *Proposed real-time DDOS mitigation system* section eliminates the above complexities. No DNS changes are required: the operator retains full control over its DNS infrastructure. No traffic rerouting occurs: legitimate traffic flows directly from source to destination through normal BGP paths, and only malicious traffic is dropped at the ISP/IXP/gateway level before it reaches the victim's link. No privacy exposure arises: traffic is simply dropped (not redirected, inspected, and forwarded), so no user data passes through third-party infrastructure and no deep packet inspection is needed. The model is scalable at national level with zero marginal infrastructure cost, because DNSC manages a centralised whitelist/blacklist that is propagated to participating ISPs and IXP operators. Whether

Table 1. *Comparative analysis: commercial scrubbing services vs. proposed DNSC-coordinated upstream filtering model*

Criterion	Commercial scrubbing	Proposed DNSC model
DNS changes	Required (delegation to provider)	None
Traffic rerouting	All traffic via scrubbing centre	None – only malicious traffic dropped upstream
Privacy exposure	Full traffic inspection by third party	None – binary drop/forward, no DPI
SSL/TLS handling	Key sharing or proxy termination	Not applicable
Bandwidth overhead	Doubled (to/from scrubbing centre)	Zero – malicious traffic never reaches victim
Cost per entity	≥ 700 EUR/month	Zero marginal cost
Scalability	Linear cost increase per entity	Constant – filtering at shared upstream nodes
GDPR compliance	Requires DPA with foreign processor	Inherently compliant – no data processing
Vendor lock-in	High (DNS, BGP, tunnel dependencies)	None – uses standard nftables/BGP
National backbone load	Increases during attacks	Decreases during attacks

one micro data centre or one thousand are under simultaneous attack, the filtering rules are applied at the same upstream nodes. Table 1 summarises the comparative analysis between current commercial scrubbing services and the proposed DNSC-coordinated upstream filtering model.

OPERATIONAL MECHANISMS FOR INCIDENT REPORTING AND MANAGEMENT

The effectiveness of the National Cyber Security Strategy depends operationally on the speed with which incidents are detected and communicated to the competent authority. In accordance with the provisions of Directive (EU) 2022/2555 (NIS2) and the national transposition framework, the reporting procedure requires a three-stage mechanism (European Parliament & Council, 2022): an early warning within maximum 24 hours from the moment the entity becomes aware of a significant incident, signalling the existence of a potential crisis and indicating, where applicable, a cross-border attack vector; an incident notification within maximum 72 hours, updating the initial information with a preliminary severity assessment, impact on the availability, confidentiality, and integrity of data, and identified Indicators of Compromise (IoC); and a final report within maximum one month, comprising a detailed incident description, root-cause analysis, applied mitigation measures, and final cross-border impact assessment.

DNSC provides the national dedicated platform (PING – National Platform for Cyber Security Incident Management) for standardised and secured reporting. Additionally, the national unique emergency number for cyber security incidents – 1911 – is operational for critical incidents requiring immediate assistance and for reports from citizens and SMEs not necessarily falling under NIS2 mandatory reporting obligations (DNSC, 2024). Analysis of statistical data from 2024–2025 reveals a direct correlation between the use of automated reporting tools and the reduction of mean incident response time. However, the major challenge remains the capacity of targeted entities (especially from the public sector and healthcare) to collect the technical data necessary for the 72-hour report under conditions where ransomware or DDoS attacks may block access to their own monitoring and logging systems (ENISA, 2024).

PROPOSED REAL-TIME DDOS MITIGATION SYSTEM

Problem statement and design rationale

As demonstrated in the *Comparative analysis of DDOS attacks: Global vs. National Context* section, the majority of DDoS attacks are short-lived (under 30 minutes), during which the router detects the attack but is incapable of stopping it because either the attack overwhelms the available 10 Gbps bandwidth – a typical link capacity for a micro data centre in Romania in 2025 – or it exhausts the router's CPU resources. Malicious DDoS traffic, by source, falls into two categories: national and international, with the majority originating from international sources. Romania has several well-known and controlled physical entry points for international Internet traffic, and nationally, traffic capable of producing such attacks transits through Internet Exchange Points (IXPs) such as InterLAN, RONIX, and NetIX (PeeringDB, 2025). The DNSC reporting timelines of 24 hours for early warning and 72 hours for incident notification, while improved from previous frameworks, are insufficient for resolving DDoS incidents whose duration is measured in minutes rather than days.

Proposed architecture

The proposal involves a real-time DDoS mitigation system structured in five steps. In the first step (detection), the edge router – running any OS – detects the attack and analyses the implicated IP addresses, ports and attack type using flow-based analysis (sFlow, netflow) or other available solutions combined with rate-based anomaly detection. In the second step (real-time reporting), the router automatically reports the attack to DNSC via a standardised REST API; the core payload consists of a three-tuple per identified attack flow (source IP address or CIDR range, destination IP address, destination port), together with metadata such as attack-vector classification, packet-rate metrics and protocol type (but also other parameters required for the nftables rules / rate limiting). In the third step (DNSC validation), DNSC analyses the request in real time and validates that (a) the request originates from the verified owner of the destination IP address, (b) the three-tuple is correctly formulated and the destination IP belongs to a registered Romanian autonomous system, and (c) upon validation, the approved blocking rules are transmitted to all relevant ISPs and IXP operators. In the fourth step (upstream filtering), DNSC communicates the validated three-tuple blocking rules to the IXP and international gateway operators in real time; at the ISP/IXP router level these are translated into nftables set elements that selectively drop only the matching traffic flows. This is more granular than traditional BGP blackhole communities (RFC 7999, 2016; Kumari & McPherson, 2009), which null-route entire IP prefixes, because the three-tuple approach preserves legitimate traffic to and from both the source and destination addresses on non-targeted ports. In the fifth step (traffic blocking), the IXP and international gateways drop packets matching the DNSC-validated three-tuples; rules include configurable timeouts for automatic expiry, and DNSC can issue unblock commands to remove specific three-tuples when the attack subsides.

Figure 1 summarises the operations and message flows between the three categories of participants – end users, the DNSC platform, and ISP/IXP/international-gateway operators – covering both the one-time pre-registration phase and the five real-time mitigation steps described above. It also illustrates the confirmation and debug-reporting feedback loops detailed within the *ISP and gateway pre-registration* and *Extended rule schema, debug mode and emergency escalation* sections.

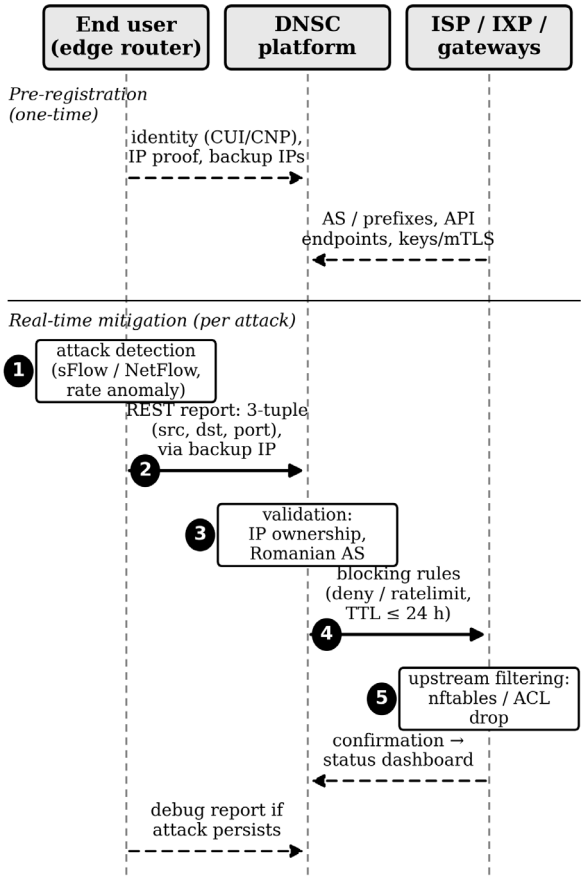


Figure 1. *Interaction diagram of the proposed DNSC-coordinated real-time DDoS mitigation system: one-time pre-registration and the five real-time mitigation steps between the involved entities*

DNSC platform architecture and communication protocol

The practical viability of the system depends on a well-defined communication protocol between three categories of participants: end users (micro data centre operators, public institutions, enterprises), ISPs/IXP gateways, and DNSC as the central coordination authority. A critical design requirement is platform agnosticism: gateway routers may run proprietary vendor firmware (Cisco IOS, Juniper Junos, MikroTik RouterOS), open-source distributions (OpenWrt, VyOS, Buildroot-based systems), or general-purpose Linux servers acting as routers. The communication protocol therefore relies on universally available mechanisms – HTTPS REST API calls and simple shell commands executable on any platform with root console access (Cisco, 2023). DNSC operates a lightweight platform comprising three core components: an API layer for receiving and dispatching blocking/unblocking requests; a web-based pre-registration portal for end users to register IP ownership; and a web-based pre-registration portal for ISPs and gateways to register as filtering endpoints.

End-user pre-registration and IP ownership verification

Before an end user can submit blocking requests, they must pre-register on the DNSC platform and prove ownership of the public IP addresses they wish to protect. This pre-registration is essential to prevent abuse: without it, a malicious actor could submit blocking requests for IP addresses they do not control, effectively weaponising the system as a denial-of-service tool itself. The end user accesses the DNSC web portal and creates an account with identity verification (e.g., tied to a Romanian business registration number – CUI – or personal CNP, consistent with existing eIDAS electronic identification frameworks). For each declared IP range, the user must provide proof of control, either through manual verification (uploading a signed letter or contract excerpt from the ISP confirming the allocation, e.g., a RIPE database assignment) or through automatic verification leveraging the RIPE Database RDAP/Whois API, a TXT record in the reverse DNS zone, or an HTTP challenge on a specified port – analogous to ACME domain validation used by Let's Encrypt for SSL certificates (ISRG, 2019). If DNSC establishes a bilateral protocol with Romanian ISPs (already required to cooperate with DNSC under NIS2 obligations), the ISP could confirm allocations via a dedicated API endpoint, enabling fully automated real-time verification.

Backup IP registration and out-of-band communication

A critical consideration is communication resilience during an active attack. If all of a user's public IP addresses are under DDoS attack, the user's router may be unable to communicate with DNSC via those addresses – the very traffic carrying the blocking request would be drowned out by the attack traffic. The pre-registration process therefore requires the user to declare one or more backup IP addresses from different upstream providers or separate network segments that are not publicly associated with the user's primary infrastructure and are therefore unlikely to be targeted simultaneously (e.g., a mobile data connection, a secondary ISP link, or a VPN exit point). These backup IPs are stored in the DNSC database and are pre-authorised to submit blocking requests on behalf of the registered primary IPs. The attack scenario then unfolds as follows: the user's primary IPs come under DDoS attack and become unreachable, the user's monitoring system (or the operator manually) initiates a blocking request from the pre-registered backup IP, DNSC validates that the backup IP is pre-registered and associated with the declared primary IPs, that the destination IPs in the request belong to the same set, and that the source IPs to be blocked are not in the user's own registered set (preventing self-inflicted blocking). Since the backup IP was pre-authorised, DNSC propagates the blocking rules to all participating ISPs/gateways within seconds.

ISP and gateway pre-registration

On the receiving side of DNSC communications, ISPs and gateway operators must also pre-register on the platform. DNSC can require the person actually pre-registering the company to be a cyber security auditor as defined in ORDER SGG 559/2021. Each ISP/gateway registers the network ranges they operate (autonomous system numbers and IP prefixes), one or more API endpoints (URLs) where DNSC will POST blocking/unblocking instructions, and authentication credentials (API keys or mutual TLS certificates) for securing the communication channel. When DNSC validates a blocking request, it looks up all ISPs and gateways through which traffic to the victim's IP prefix is routed and dispatches the three-tuple blocking rules to each registered endpoint simultaneously. The receiving system – regardless of whether it runs on Linux (nftables), Cisco (ACL), MikroTik (firewall rules), or any other platform – translates the received three-tuple into the appropriate local filtering syntax and applies it. The gateway then responds to DNSC with a confirmation (or error), enabling DNSC to maintain a real-time status dashboard of which rules have been successfully applied across the national infrastructure.

Extended rule schema, debug mode and emergency escalation

While the basic three-tuple provides the core filtering specification, a production-grade system requires additional parameters. The extended schema supports four operation types: deny (silently drops all matching packets – the default DDoS mitigation action), accept (explicitly whitelists matching traffic, overriding any deny rules), ratelimit (limits the matching traffic to a specified packets-per-second or megabits-per-second threshold), and forward (redirects matching traffic to a designated scrubbing or analysis endpoint). Each rule includes explicit temporal controls: a duration expressed in seconds, with a maximum auto-reset of 86,400 seconds (24 hours) as a safety mechanism. Regardless of the duration requested by the end user, all rules automatically expire after one day unless explicitly renewed, preventing stale blocking rules from persisting indefinitely. Each request optionally includes a report_ddos flag: when set, the blocking request

simultaneously generates a formal incident report in the DNSC incident-tracking system, automatically satisfying the NIS2 24-hour early warning requirement and reducing the administrative burden on the defender.

A particularly challenging operational scenario arises when DNSC has dispatched blocking rules to all registered ISPs and gateways, yet the attacked router continues to receive malicious traffic. The protocol therefore includes a debug-reporting mechanism: when the attacked router detects that traffic matching a previously confirmed three-tuple continues to arrive, it sends a debug report to DNSC containing the original request ID, packet captures or flow summaries showing continued matching traffic, the directly connected upstream peer's IP, and the arriving interface and VLAN data. Upon receiving such a report, DNSC cross-references the original rule against the list of ISPs/gateways that confirmed rule application, identifies any that did not respond or reported errors, and – if an ISP is found to be in breach of its obligation to apply DNSC-validated filtering rules under the NIS2 cooperation framework – initiates the appropriate regulatory escalation. This transforms the system from a purely reactive filtering tool into an accountability framework where non-compliance can be investigated with concrete evidence.

In extreme scenarios where targeted three-tuple blocking is insufficient – for example, when an attack employs rapidly rotating source IPs across thousands of subnets – the system provides emergency escalation options. Upstream rate limiting allows the user (or DNSC on the user's behalf) to request a temporary rate limit on all inbound traffic destined for a specific IP address, regardless of source. Geographic restriction allows the temporary dropping or rate-limiting of all traffic originating from outside Romania for a specific destination IP, which mitigates most attack vectors while preserving domestic connectivity, given that the majority of DDoS traffic originates from international sources. Both escalation mechanisms are subject to elevated DNSC authorisation and to a shorter maximum duration (typically 4 hours) given their broader impact on service availability.

Bidirectional reporting: ISPs as attack reporters

The architecture described thus far assumes that blocking requests originate exclusively from end users. However, ISPs and gateway operators are themselves targets of DDoS attacks and are often the first to detect large-scale attack patterns affecting multiple downstream customers simultaneously. The system therefore provides a symmetric capability where ISPs and gateways can also report attacks and request DNSC-coordinated mitigation using the same protocol. Since the ISP is a pre-registered network operator, DNSC can validate that the destination IPs fall within the ISP's announced prefixes without requiring individual end-user ownership verification, enabling ISPs to protect their customers proactively. ISP-originated reports also provide DNSC with a higher-altitude view of the attack landscape; while an individual micro data centre sees only the traffic destined for its own IPs, an ISP observes aggregate patterns across thousands of customers and can detect coordinated botnet campaigns invisible to any single victim. This bidirectional reporting architecture ensures that the DNSC platform serves as a true national coordination hub.

Strategic advantage of filtering at country entry points

A fundamental argument for this architecture is that by blocking malicious traffic at the country's entry points (international gateways and IXPs), DDoS attacks become extremely difficult to execute successfully. At these points, there exists sufficient bandwidth capacity and processing power to mitigate

the vast majority of attacks – or at minimum substantially more capacity than a micro data centre with a typical 10 Gbps link. International gateways operated by Tier-1 and Tier-2 providers in Romania (DIGI/RCS&RDS, Orange, Vodafone) maintain aggregate capacities in the hundreds of Gbps to Tbps range (DIGI, 2024; Orange, 2024), making even large-scale attacks manageable when filtered at this level. This approach extends the established Remote Triggered Blackhole (RTBH) practice (Cisco, 2023; Kumari & McPherson, 2009) with automated detection and DNSC-coordinated validation for a nationally coordinated response, and remains compatible with BGP FlowSpec where supported (RFC 8955, 2020). The protocol elements described above represent the minimum viable specification: a production implementation by DNSC would necessarily address authentication and authorisation (OAuth 2.0, mutual TLS), message signing for non-repudiation, redundancy of the platform, integration with the existing PING incident-management workflows, and legal frameworks for ISP cooperation obligations under NIS2 – details that fall outside the scope of this paper and are left for DNSC and the relevant governmental bodies to define in consultation with ISP operators and the broader cybersecurity community.

Application settings and future case studies

To ground the proposed architecture in operational reality, several concrete application settings can be identified, which are also intended to serve as the basis for future case studies. A first, retrospective scenario is the wave of hacktivist DDoS campaigns against Romanian public-administration websites observed since 2022, in which government portals and transport-sector services were rendered unavailable by relatively modest volumetric attacks; replaying the traffic profiles of such incidents in a controlled testbed would allow a direct measurement of the time-to-mitigation achievable with DNSC-validated three-tuple filtering compared with the manual, per-institution response actually applied. A second setting is the regional micro data centre or hosting provider operating on a typical 10 Gbps uplink, for which even a mid-sized attack of 20–50 Gbps is saturating; this is the archetypal beneficiary of upstream filtering, since no on-premises equipment can help once the last-mile link is full. A third setting covers public healthcare and educational institutions – essential or important entities under NIS2 that rarely have the budget for commercial scrubbing subscriptions – where availability losses translate directly into interrupted patient services or blocked admission and examination sessions. Finally, election periods and other moments of heightened geopolitical tension constitute a cross-cutting setting in which coordinated campaigns target many public entities simultaneously, precisely the situation in which per-victim commercial mitigation scales worst and centrally coordinated filtering scales best.

As future work, the authors propose a staged validation path: (i) a laboratory testbed reproducing the end-user / DNSC / ISP roles on virtualised Linux routers, measuring rule-propagation latency and end-to-end time-to-mitigation for the attack vectors documented in the *Comparative analysis of DDOS attacks: Global vs. National Context* section; (ii) a pilot deployment with a Romanian IXP (e.g., InterLAN) and a small group of volunteer operators; and (iii) a comparative case study running the same recorded attack traces against both the proposed mechanism and a commercial scrubbing service, quantifying the differences in mitigation delay, collateral damage to legitimate traffic, and cost.

CONCLUSIONS AND RECOMMENDATIONS

Corroborating the statistical data presented, the conclusion emerges that Romania presents a vulnerable attack surface in the face of the new generation of hyper-

volumetric threats. The current infrastructure and measures provided in the 2022–2027 Strategy require urgent recalibration. The proposed real-time mitigation system offers several decisive advantages over the current approach: it reduces response time from hours and days to seconds; it leverages existing infrastructure without requiring massive capital investment; it creates a nationally coordinated defence where DNSC validates and orchestrates filtering rather than each entity defending independently; and, by filtering at country-level entry points where bandwidth capacity is orders of magnitude greater than at individual endpoints, even large-scale attacks become manageable.

The proposed DNSC platform architecture – comprising a REST API, end-user pre-registration with IP ownership verification and backup IP provisioning, and ISP/gateway pre-registration – demonstrates that the coordination layer can be implemented with standard web technologies at minimal cost. The extended rule schema supports multiple operation types (deny, accept, ratelimit, forward), time-bounded rules with a 24-hour auto-reset safety mechanism, optional NIS2 incident reporting, debug diagnostics for identifying non-compliant transit paths, and emergency escalation options including upstream rate limiting and geographic traffic restriction. The bidirectional reporting capability enables ISPs to participate not only as rule enforcers but also as attack reporters, creating a comprehensive national threat-intelligence picture. Integration requires nothing more than a single shell script per participant, executable on any device with root console access, making the system deployable across the full spectrum of network equipment from embedded routers to carrier-grade gateways with virtually no modifications to existing infrastructure.

Compared to commercial solutions such as those offered by DIGI Romania (from 700 EUR/month) or international providers like Cloudflare and Akamai (from hundreds to thousands of EUR monthly), the proposed system provides a complementary, infrastructure-level protection layer at negligible marginal cost. It is essential to implement proactive protection mechanisms capable of mitigating not only volumetric bandwidth attacks but also high-intensity processing-resource attacks, exacerbated by the offensive use of Artificial Intelligence. The authors recommend that DNSC and the relevant governmental bodies urgently evaluate the implementation of the proposed automated upstream filtering system as a critical component of Romania's updated cyber security posture.

REFERENCE LIST

Apostu, A., Gheorghe, S., Hîji, A., Cleju, N., Pătrașcu, A., Rusu, C., Ionescu, R. & Irofti, P. (2025) Detecting and mitigating DDoS attacks with AI: A survey. arXiv preprint arXiv:2503.17867. Available at: https://arxiv.org/abs/2503.17867 [Accessed 28 August 2026].

Cisco (2023) *Remotely triggered black hole filtering – destination based and source based.* Cisco Systems Technical Documentation. Available at: https://www.cisco.com/c/dam/en_us/about/security/intelligence/blackhole.pdf [Accessed 15 March 2026].

Cloudflare (2025) *DDoS threat report for 2024 Q4.* Cloudflare Radar. Available at: https://radar.cloudflare.com/reports/ddos-2024-q4 [Accessed 10 March 2026].

DedicatServer (2024) *Company profile – DDoS protection infrastructure in Romania.* Available at: https://dedicatserver.ro/en/company/ [Accessed 20 March 2026].

DIGI (2024) *DIGI Communications N.V. annual report 2024.* Bucharest, DIGI Communications N.V.

DIGI (2025) *Protecție Anti-DDoS – DIGI Romania S.A.* Available at: https://www.digi.ro/campanii/protectie-anti-ddos [Accessed 25 March 2026].

DNSC (2024) *Raport anual DNSC 2023: Starea securității cibernetice în România.* București, Directoratul Național de Securitate Cibernetică.

DNSC (2025) *Raport de activitate DNSC 2024: Evenimente și incidente de securitate cibernetică monitorizate.* București, Directoratul Național de Securitate Cibernetică.

ENISA (2024) *ENISA threat landscape 2024.* Athens, European Union Agency for Cybersecurity. Available at: https://www.enisa.europa.eu/publications/enisa-threat-landscape-2024 [Accessed 12 March 2026].

European Parliament & Council (2016) *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data (General Data Protection Regulation – GDPR).* Official Journal of the European Union, L 119, pp. 1–88.

European Parliament & Council (2022) *Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union (NIS2).* Official Journal of the European Union, L 333, pp. 80–152.

Government of Romania (2022) *Strategia de Securitate Cibernetică a României pentru perioada 2022–2027.* Monitorul Oficial, nr. 2 bis, 3 ianuarie 2022.

GSL (2024) *Record-breaking DDoS attack: 3.15 billion packets per second.* Global Secure Layer Blog, August 2024. Available at: https://globalsecurelayer.com/blog/ [Accessed 15 March 2026].

ISRG (2019) *ACME protocol – automatic certificate management environment. RFC 8555.* Internet Security Research Group / Internet Engineering Task Force. Available at: https://www.rfc-editor.org/rfc/rfc8555 [Accessed 20 March 2026].

Kumari, V. & McPherson, D. (2009) *Remote triggered black hole filtering with unicast reverse path forwarding (uRPF). RFC 5635.* Internet Engineering Task Force.

Microsoft (2024) *Microsoft digital defense report 2024.* Redmond, WA, Microsoft Corporation. Available at: https://www.microsoft.com/en-us/security/security-insider/microsoft-digital-defense-report-2024 [Accessed 12 March 2026].

NETSCOUT (2025) *NETSCOUT DDoS threat intelligence report – 1H 2025.* Westford, MA, NETSCOUT Systems, Inc. Available at: https://www.netscout.com/threatreport [Accessed 18 March 2026].

Orange (2024) *Orange Business annual security report 2024 – Romania.* Paris, Orange S.A.

PeeringDB (2025) *Romanian Internet Exchange Points.* Available at: https://www.peeringdb.com/ [Accessed 20 March 2026].

Radware (2024) *DDoS protection pricing guide.* Radware Cyberpedia. Available at: https://www.radware.com/cyberpedia/ddos-protection/ddos-pricing/ [Accessed 22 March 2026].

RFC 7999 (2016) *BLACKHOLE community.* Internet Engineering Task Force. Available at: https://www.rfc-editor.org/rfc/rfc7999 [Accessed 15 March 2026].

RFC 8955 (2020) *Dissemination of flow specification rules.* Internet Engineering Task Force. Available at: https://www.rfc-editor.org/rfc/rfc8955 [Accessed 15 March 2026].

Vevera, A.-V., Crăciun, A. C., Dumitrache, M., Sandu, I., Rotună, C.-I. & Bostan, R. A. (2025) Cybersecurity challenges in managing domain names. From DNS to ENS in the Web3 era. Romanian Cyber Security Journal. 7(1), 97-112. doi: 10.54851/v7i1y202508.



This is an open access article distributed under the terms and conditions of the
Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License.