When DDoS Protection Starts Hurting Operations

Building an Automation Layer for Anomaly Correlation and Analyst Decision Support

Nishal Rai
Team Lead (Application Security Enginer) |Digital Network Solution
sudo@nishalrai.com.np
Web: nishalrai.com.np

Before Anything: yes, I run this

The network I operate

- F5 Silver-Partner protecting telecom & ISP infrastructure in Nepal
- Out-of-band DDoS detection on distributed, bursty carrier traffic
- GenieATM collector/controller + F5 AFM scrubbing, in production
- My team handles the anomaly events on shift — I'm not the vendor

$ **whoami**
> Graduated from Islington College in 2021 with a **B.Sc. (Hons) Computer Networking and IT Security**
> And, currently working as Team Lead of the Application Security department at Digital Network Solution Ltd., Naxal, Kathmandu

$ **stack**
> GenieATM, n8n, OpenSearch, F5 AFM, Third Party security intel platform

$ **telemetry**
> sFlow 1:1000 + SNMP @ edge

Detection ≠ Decision

Detection without correlation is **operational noise**.

15-30 min

Old manual triage per event — login, copy event ID, check 3 reputation sources × N IPs, judge, document

~90 sec

Same verdict, enriched & correlated, delivered to the analyst by the pipeline

every FP

A false positive isn't an annoyance—it's an unnecessary scrub, and that's an SLA event

*FP - False Positive

When does protection start hurting?

WHY When every detection becomes a manual investigation

Anomalies fire faster than analysts can qualify them. The platform detects; the human still has to find the event, pull reputation, and judge — at scale, detection volume becomes a queue, not an answer.

HOW When automatic baselines meet bursty ISP reality

Traffic varies by region, time, CDN and upstream. Intermittent spikes — a Dashain banking surge, an IPO window, a viral cache-fill — trip false detection. Bursty patterns ≠ attacks, but the baseline can't tell.

WHEN When event volume peaks during normal business surges

The worst false positives arrive exactly when legitimate traffic is highest — month-end salary processing, festival load — so the noise is loudest when you can least afford a wrong call.

WHERE When mitigation itself becomes the incident

Every unnecessary F5 AFM redirect adds latency, stresses BGP convergence on core routers, and flaps legitimate client sessions. Being 'protected' starts degrading the very traffic you protect.

Full Architecture: Detection to Decision Pipeline

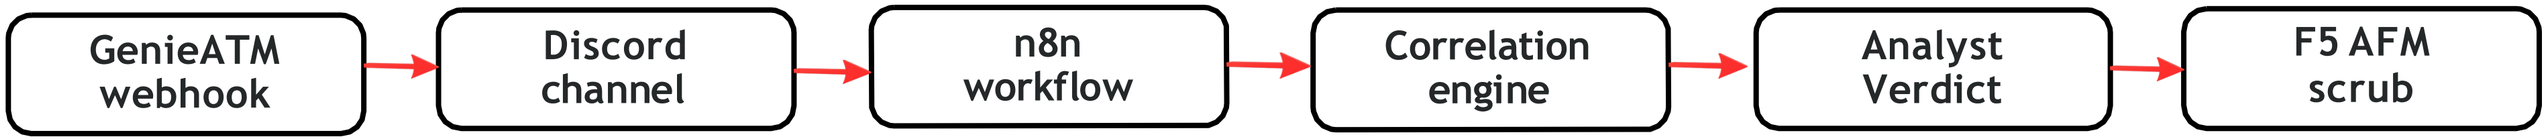

⛔ **NO EXTERNAL AI → NO DATA LEAVES THE NETWORK**

AbuseIP DB

GreyNoise

IPinfo

Only per-IP lookups leave. Event data, scoring, and history stay on-prem (OpenSearch on RHEL 9.x).

Design rationale — self-hosted on purpose

n8n (event-driven, self-hosted)
OpenSearch (no licence friction, long-term storage)
Discord (reliable webhook intake + channel separation).
Lightweight reputation only — no model phoning home.

Automation in action — what happens in 90 seconds?

01

GenieATM Fires webhook

When anomaly ddos event occurs

02

T+2s Discord receives

Within a second, Discord channels receive the raw event

03

T +5s n8n pattern

n8n pattern check passes

04

GenieATM Call

T +8s API call to GenieATM for full event detail and parse them

05

T +12s IP enrichment

IP parallel enrichment (AbuseIPD + GreyuNoise + IPInfo simulatenously

06

T +60-90s Report

Process the information and build the correlated report for the summary

07

Report to Discord

The correlated report with the summary and verdict is send to the Discord Channel

08

Action: ignore or mitigate

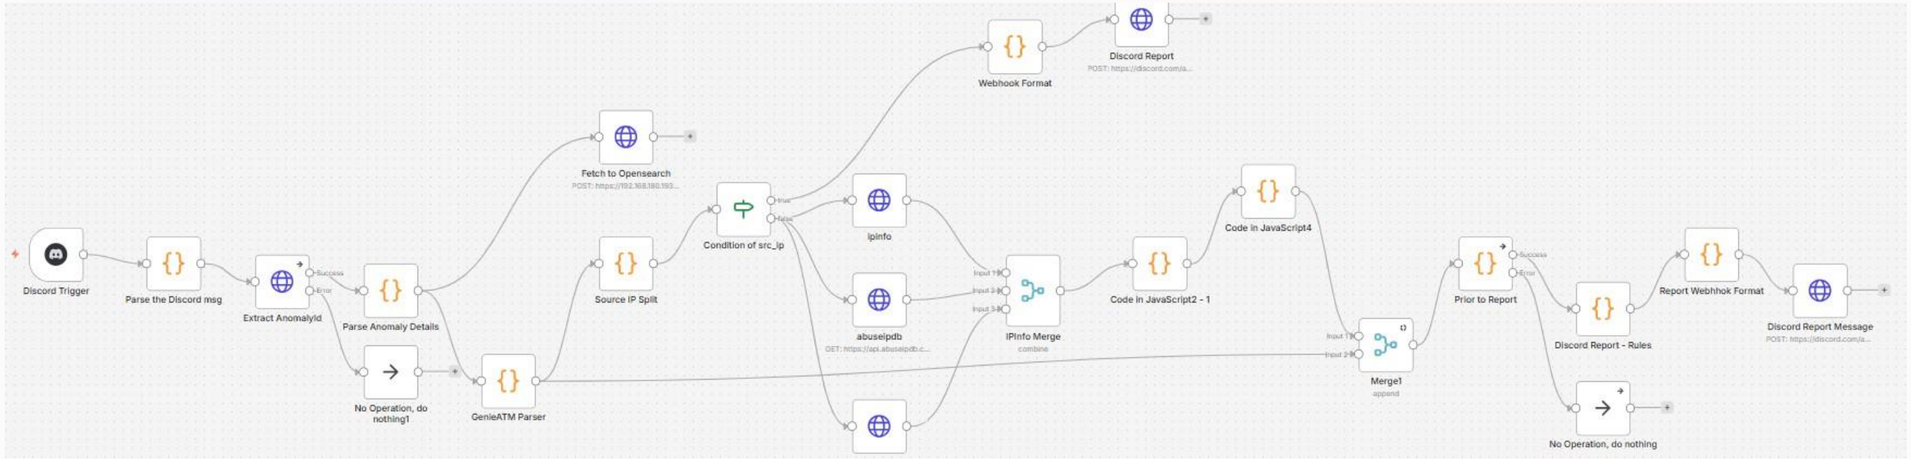

How is a verdict decided?

Dual scoring engine (per IP)
Each IP gets a threat score (AbuseIPDB weighted by report count, recency decay, DDoS-category relevance; GreyNoise malicious with staleness penalty) AND an independent FP score (RIOT safe-tags, benign classification, mail/gov/edu hostnames, CDN membership). Hard override: cloud-provider ASNs are blocked from safe-trust — that's where rented botnets live.

per-IP verdict
Confirmed Threat / Suspicious / Likely FP / Mixed Signals / Data Unavailable. No IP above threat-score 60 can be 'likely FP' regardless of how much FP credit it stacks.

volume-weighted aggregation
IP scores are aggregated and weighted by traffic share. An IP pushing 80% of the bandwidth matters far more than one contributing 2%.

Behavioural Analysis
Do the TCP flags match a real flood or just normal ACK traffic? Is the threshold exceedance meaningful or barely 1.5×?
Did it last 8 seconds or 8 minutes?

Composite, Probability & Recommendation
Signals merge into a composite score, pushed through a sigmoid to an attack probability 0 – 1; confidence drops if sources fail or sample is thin. Output: auto-mitigate · escalate · suppress · pull PCAP.

The analyst at 3am

The most relatable pain in the room: alert fatigue. Here's the before/after of a real shift.

Before — manual workflow (15-30 min/event)

1. Alert fires — analyst paged
2. Login to GenieATM, copy anomaly event ID, search (2-3 min)
3. Check each source IP in AbuseIPDB + GreyNoise + ASN lookup (≈5 min × N IPs)
4. Judgment call → escalate or dismiss
5. Document, build report, close case

Multiply by every event, every shift.
That's the fatigue — and the missed events that fatigue causes.

After — enriched embed replaces the runbook

- The 3am runbook is now executed by the pipeline, not by a tired human
- Analyst reads a structured verdict in ~90 seconds, then decides
- OpenSearch holds full event history → clean shift handover
- One analyst now handles 5× the event volume

Research → done by pipeline
Decision → Stays human

RTBH blast radius vs FlowSpec precision

At the peering/transit layer, the mitigation choice is itself an operational risk decision.

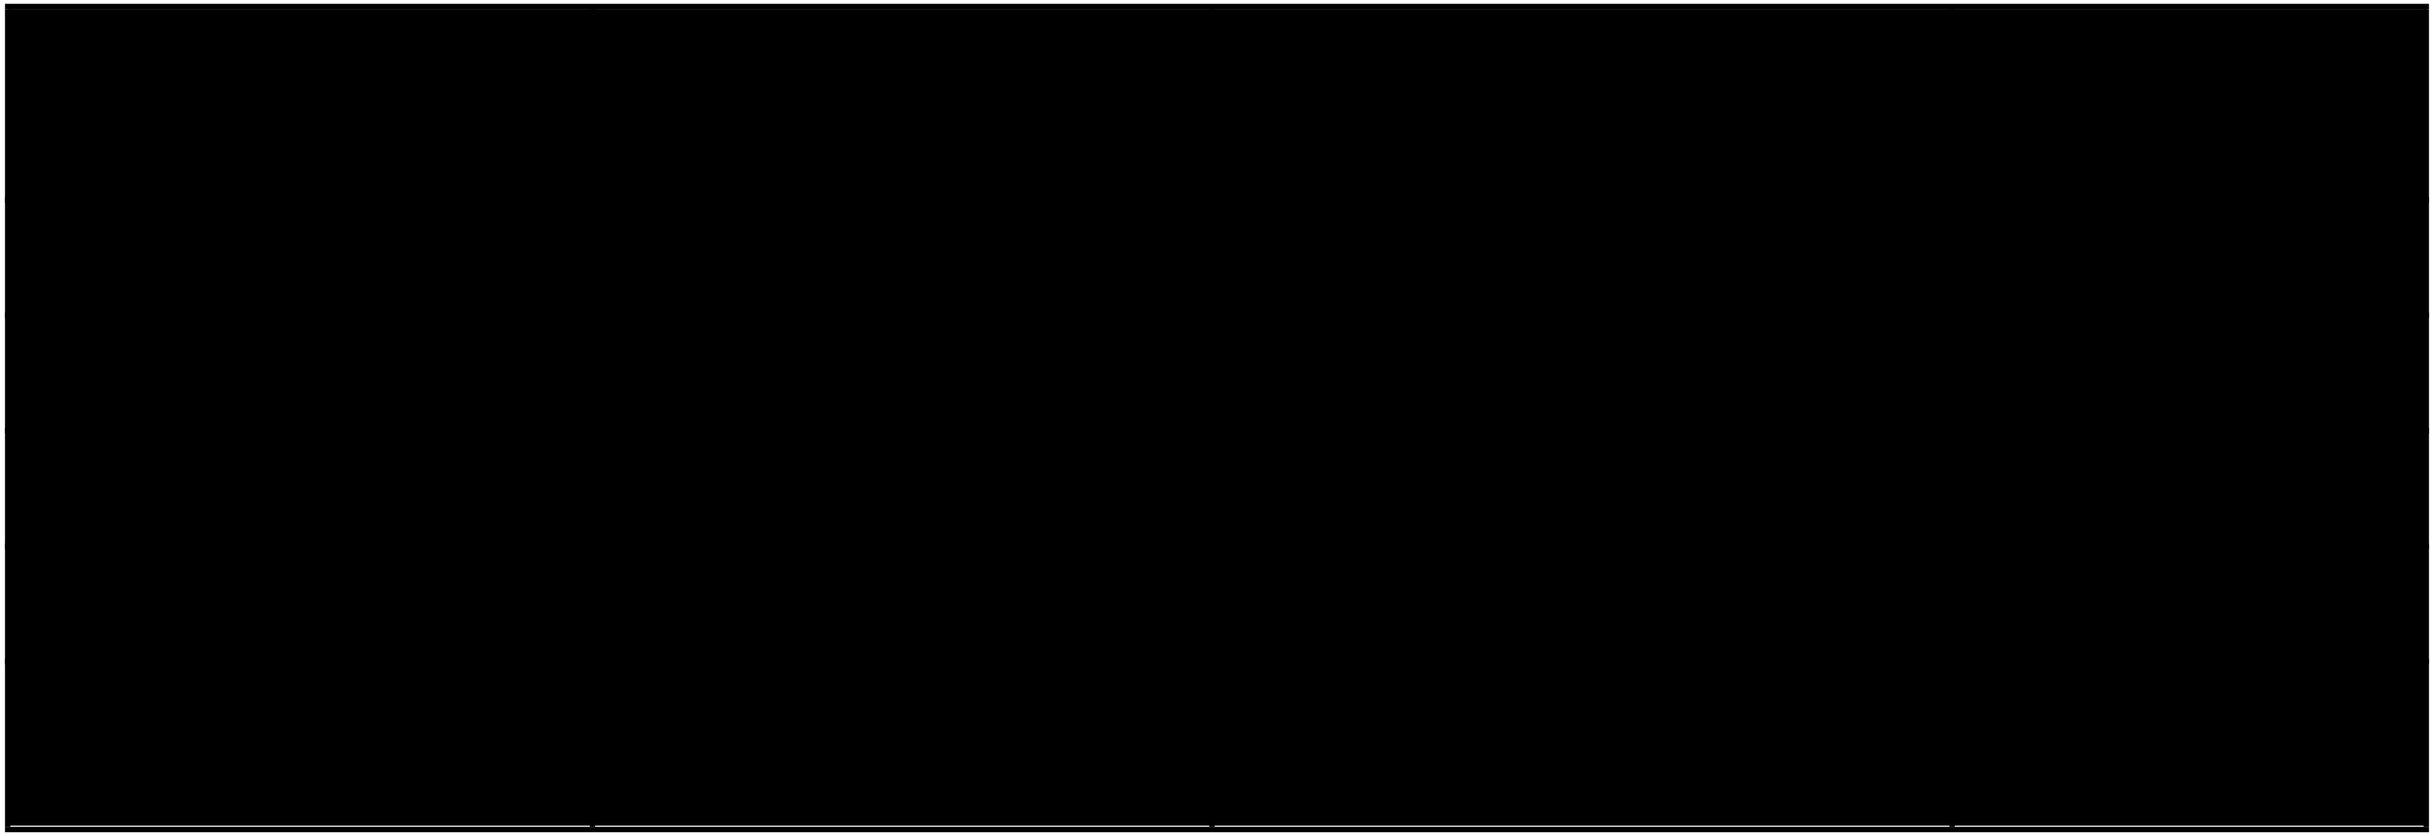

The operator truth

RTBH completes the attacker's job at the edge, the destination goes dark. FlowSpec is on our roadmap precisely so we can drop attack traffic at ingress without blackholing the customer. RTBH is the last resort, not the first.

Every false redirect is an SLA event

+RTT

Clean traffic steered through F5 AFM adds latency for every session in the protected zone

BGP churn

Each redirect stresses convergence on core routers; oscillating detection causes state flapping

SLA Hit

Wholesale / enterprise customers feel it where a "protection" action becomes a service-impact event

So the metric that matters

Cutting unnecessary scrubbing by 50–60% isn't an efficiency stats, it's fewer SLA events, less BGP convergence stress, and stable client sessions.
That's the operator framing.

The one that didn't get scrubbed

Event walk-through [A381072]

1. Event: May 2 2026, burst on a host, GenieATM labels it "Host TCP Traffic"
2. Detection said: 4,200 Kbps vs 1,450 threshold → 2.9× multiplier, 115s, severity YELLOW
3. Pipeline added: top source = 103.140.53.92, 99% volume → a Cisco Email Security Appliance, Nepal Police (government)
4. Behaviour: 99.43% ACK / 0.57% SYN → established sessions, not a flood. Flag mismatch flagged.
5. Outcome: F5 AFM scrub NOT triggered.Government email kept flowing.

DDoS Event Assessment — A381072

Verdict	LIKELY FALSE POSITIVE
Attack probability	9.4%
Confidence	65%
IP Net Score	−34.7
Composite	−56.7
Dominant source	Nepal Police · gov
Recommendation	REVIEW THRESHOLD

A measurable avoided service impact on government infrastructure, the exact event that would've triggered an unnecessary scrub before the pipeline existed.

A real flood, scrubbed in 90 seconds

Event walk-through

1. Event: volumetric UDP/SYN flood on a transit-facing prefix
2. Detection said: high-multiplier bps/pps spike, multiple high-abuse source IPs across diverse ASNs
3. Pipeline added: dominant sources = high AbuseIPDB confidence + GreyNoise malicious, cloud-ASN safe-trust blocked
4. Behaviour: flag distribution + duration + amplification ports consistent with a real flood
5. Outcome: F5 AFM redirect confirmed; clean traffic reinjected, malicious dropped

DDoS Event Assessment — confirmed

Verdict	LIKELY REAL ATTACK
Attack probability	> 0.75
Action	AUTO-MITIGATE / ESCALATE
Time to mitigation	~90 s

90s

vs 15-30 min manual — a 10-20× faster response

3×

better F5 AFM utilisation — capacity reserved for confirmed danger

State flapping, eliminated

Before

- Oscillating detection redirected legitimate client sessions to the scrubber and back
- Sessions reset - users saw intermittent failures during "protection."
- Repeated redirects stressed BGP convergence on core routers

After [self-improving loop]

- Confirmed events -> tighten GenieATM thresholds + refine F5 AFM policy per attack pattern
- Flapping for legitimate sessions eliminated; noise down 50-60%
- The loop hardens the whole stack incrementally with every incident handled

The pipeline feeds the infrastructure, not just the analyst

Each verdict makes the next detection better. That's the difference between a dashboard and an operational system.

Small attacks, outsized damage

Documented Nepal incidents

- Jan 2023 — DDoS took ~1,500 govt sites offline; immigration paralysed for about 4 hours
- Jan 2024 — GIDC attack: 400+ govt sites down, Tribhuvan Airport immigration disrupted; foreign traffic geo-blocked as emergency response
- Mar 2025 — 400+ sites down again + ShadowLeak breach of PM-office data (~100k rows offered ~$1,000)
- Cybercrime incidents nearly tripled 2019-2023; no Data Protection Law yet

Global scale, for contrast

- 47.1M → DDoS attacks Cloudflare mitigated in 2025 (+121% YoY)
- 31.4 Tbps → record single attack, late 2025
- 7.1 Gbps → peak for Nepali wired carriers (Netscout 2H2023)
- #1 → Telecoms = most-attacked industry; 7 of top-10

Anatomy of a correlated report

DDoS Event Assessment — A381072

- Verdict — LIKELY FALSE POSITIVE
- Traffic — 4200 Kbps · thr 1450 · 2.9× · 400 pps
- Source intel — 2 IPs · ASN diversity 2
- Behaviour — 99.43% ACK — not a flood
- Scoring — net −34.7 · prob 9.4% · conf 65%

Event Overview

GenieATM API: event ID, severity, status, target, protocol, ports, window.

Traffic metrics

GenieATM anomaly data: detected vs threshold bps, multiplier, pps, avg pps/source.

Source Intelligence

AbuseIPDB + GreyNoise per IP; ASN diversity and geo - IPinfo.

Threat Assesstment

correlation engine: dual score, volume-weight, behavior, and sigmoid probability.

Anatomy of a correlated analyst report [Report Sample]

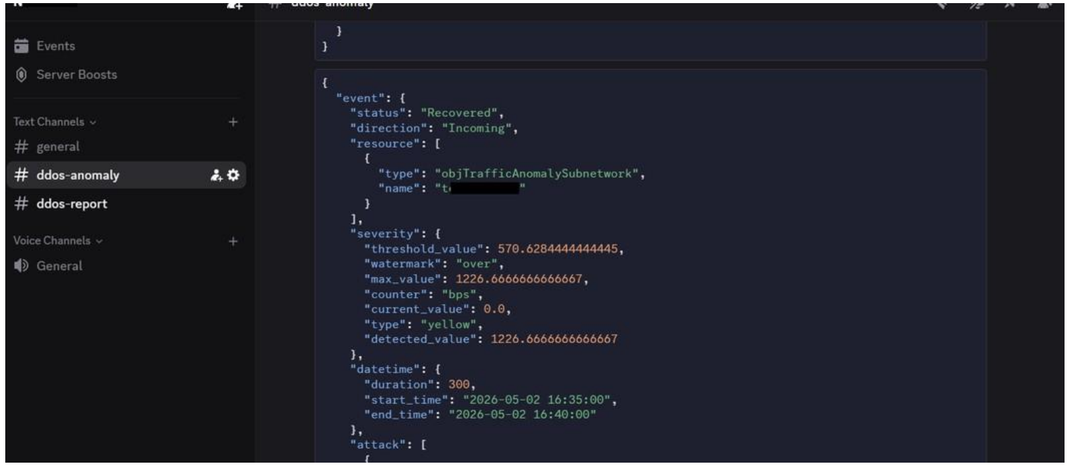

Where the tool stops

Signalling-plane DDoS is out of scope

Our sFlow/SNMP telemetry sees IP-layer volumetric and L7 floods — not control-plane attacks against SS7/Diameter (4G) or HTTP/2 SBI & GTP-C (5G). Mobile operators in this room: that's a different surface. Consult GSMA FS.11 (SS7), FS.19 (Diameter), FS.20 (GTP).

External enrichment is a dependency

AbuseIPDB / GreyNoise / IPinfo are third-party APIs. They rate-limit and they go down. The scoring engine degrades gracefully (confidence drops, MIXED SIGNALS)—but the dependency is real. Free tiers also cap daily lookups; high event volume needs caching or paid tiers.

No HA on n8n yet

If the pipeline dies, we fall back to the manual workflow - GenieATM keeps detecting, and F5 AFM keeps scrubbing on native thresholds. But the n8n layer isn't highly available today. Known gap, on the list.

What you actually need

- Any webhook/API-capable detector—not just GenieATM: FastNetMon, Wanguard, or a custom sFlow collector all work
- n8n (self-hosted) for the workflow · OpenSearch for storage/dashboards
- Free-tier enrichment: AbuseIPDB · GreyNoise · IPinfo
- Your existing scrubber: Radware DefenseProPro, A10 Thunder TPS, F5 AFM, or open-source Gatekeeper

Future Enhancements: The Next Layer

Whats get built next - beyond the current production Scoring

Whats get built next - beyond the current production Scoring

FlowSpec – Beyond Automated Mitigation

BGP FlowSpec rules propagated to edge upon confirmed attack. Upstream signaling enables traffic dropping at ingress before it reaches the F5 AFM scrubber, reducing scrubbing load for volumetric events.

NEXT

ML-Based Confidence Scoring

Replace rule-weighted scoring with a trained lightweight model. Train on historical OpenSearch event data (labeled FP vs TP). Output: probabilistic confidence score per event, not static thresholds.

NEXT

Cross-ISP Threat Intelligence Federation

Share enriched anomaly IOCs (source ASNs, confirmed attack prefixes) with regional operators via MISP or Shadowserver feeds. SANOG/npNOG community participation model for shared DDoS visibility.

NEXT

“ Any Nepali ISP or npNOG member can run this stack. ”

Key Takeaways

01

Detection is only half the work
GenieATM detects. The automation pipeline decides. F5 AFM scrubs. Each component has a distinct role – and the gap between detection and scrubbing decision is where operational pain lives.

02

Every false F5 AFM redirect is a service impact event
Unnecessary scrubbing introduces latency to legitimate traffic in the protected zone. False positive reduction is not just analyst efficiency – it is direct service quality protection.

03

The pipeline feeds the infrastructure, not just the analyst.
Confirmed events tighten GenieATM thresholds. Attack patterns refine F5 AFM policies. The automation loop hardens the entire stack incrementally with every incident handled.

04

Open-source, self-hosted, reproducible.
n8n, OpenSearch, free enrichment APIs. No vendor lock-in. Any Nepali ISP or SANOG/npNOG member can deploy this stack. Cross-ISP adoption creates a regional DDoS response capability.

References & Resources

Regional & Community

- APNIC - apnic.net
- Shadowserver - shadowserver.org
- Spamhaus - spamhaus.org
- npNOG - npnog.org.np

Standards & Protocols

- RFC 5635 - RTBH Filtering
- RFC 8955 - BGP FlowSpec
- RFC 7011 - IPFIX | sFlow v5 - sflow.org
- RFC 6811 - BGP Prefix Origin Validation (RPKI)

Tools

- n8n - docs.n8n.io
- Opensearch - opensearch.org
- GenieATM - genie-networks.com
- F5 AFM - f5.com
- FastNetMon - Wanguard
- pmacct · FRRs

Threat Intelligence & Enrichment

- AbuseIPDB | GreyNoise | IPinfo | Shodan | InternetDB
- bgp.tools Bulk ASN API – bgp.tools/kb/api
- Team Cymru IP-to-ASN – origin.asn.cymru.com
- Cloudflare Radar – radar.cloudflare.com

THANK YOU

Automate the enrichment. Keep the decision human.
The automation layer makes the analyst faster — not redundant.