SOLUTION BRIEF NetFlow Optimizer™ + Splunk Integration

Your Splunk ingestion bill is too high.

Raw NetFlow, sFlow, and IPFIX overwhelm Splunk with duplicates, micro-flows, and context-free IP addresses. NetFlow Optimizer sits before Splunk and fixes all three — delivering enriched, deduplicated network intelligence **at 80–90% less volume. Same visibility. Dramatically lower cost.**

80–90%	300K+	< 1 hr
Typical Splunk ingest reduction	Flows/sec processed, per instance	From install to live Splunk data

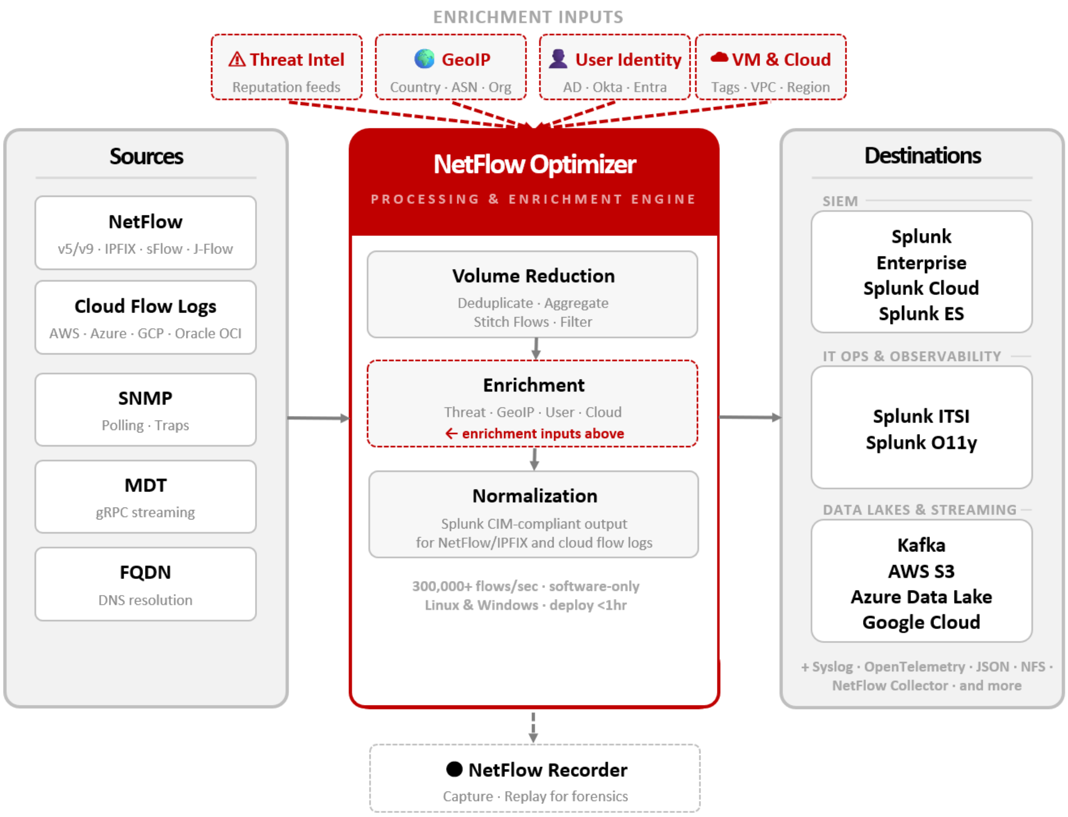

ENRICHMENT INPUTS

⚠ Threat Intel — Reputation feeds

GeoIP — Country · ASN · Org

User Identity — AD · Okta · Entra

VM & Cloud — Tags · VPC · Region

Sources

- **NetFlow** — v5/v9 · IPFIX · sFlow · J-Flow
- **Cloud Flow Logs** — AWS · Azure · GCP · Oracle OCI
- **SNMP** — Polling · Traps
- **MDT** — gRPC streaming
- **FQDN** — DNS resolution

NetFlow Optimizer

PROCESSING & ENRICHMENT ENGINE

- **Volume Reduction** — Deduplicate · Aggregate · Stitch Flows · Filter
- **Enrichment** — Threat · GeoIP · User · Cloud — ← enrichment inputs above
- **Normalization** — Splunk CIM-compliant output for NetFlow/IPFIX and cloud flow logs

300,000+ flows/sec · software-only
Linux & Windows · deploy <1hr

Destinations

SIEM

- Splunk Enterprise
- Splunk Cloud
- Splunk ES

IT OPS & OBSERVABILITY

- Splunk ITSI
- Splunk O11y

DATA LAKES & STREAMING

- Kafka
- AWS S3
- Azure Data Lake
- Google Cloud

+ Syslog · OpenTelemetry · JSON · NFS · NetFlow Collector · and more

● **NetFlow Recorder** — Capture · Replay for forensics

The Three Problems NFO Solves for Splunk

01 COST

Ingest Cost Reduction

Raw NetFlow generates millions of events per hour — many duplicated across collection points. NFO deduplicates, aggregates, and stitches flows before they touch Splunk licensing.

- → Deduplication removes cross-collector duplicates
- → Aggregation collapses micro-flows into single records
- → Flow stitching adds 50% additional reduction
- → Top N filtering focuses Splunk on high-value traffic

02 CONTEXT

Enriched Data, Not Naked IPs

Raw flows show IP addresses. NFO adds the context Splunk needs for detection, investigation, and AI/ML — before data lands in your index.

- → User identity via AD, Okta, Microsoft Entra ID
- → Threat intelligence & reputation scoring
- → GeoIP, ASN, and location mapping
- → VM names and cloud instance metadata

03 DATA NORMALIZATION

CIM-Compliant from Day One

NFO normalizes every flow source — NetFlow, IPFIX, sFlow, J-Flow, and cloud flow logs from AWS, Azure, GCP, and Oracle OCI — into a single consistent, **CIM-compliant** stream before it ever reaches Splunk. The TA-netflow handles correct field mapping on ingestion, but the normalization work is already done inside NFO.

- → All sources normalized inside NFO — before Splunk ingestion
- → Splunk HEC delivery with JSON formatting
- → Cloud flow logs (AWS, Azure, GCP, OCI) normalized alongside on-prem NetFlow

What's Included — Splunk-Specific Features

NetFlow & SNMP Analytics App

A free, pre-built app on Splunkbase with ready-to-use dashboards for network traffic analysis, firewall monitoring, cloud visibility, and SNMP device health. No dashboard configuration required.

Technology Add-On (TA-netflow)

Ensures NetFlow data lands in Splunk with correct CIM-compliant field names and sourcetype. Normalization happens inside NFO — the TA ensures Splunk indexes it correctly and makes it immediately compatible with all Splunk apps.

Zero-Touch SNMP Device Monitoring

Automatically discovers, classifies, and polls SNMP devices — sending device health KPIs (CPU, memory, interface stats, traps) directly to Splunk alongside flow data, in a single pipeline.

Content Pack for Splunk ITSI

Extends NFO data into Splunk IT Service Intelligence — adding service health views, glass tables, and KPI monitoring driven by enriched NetFlow and SNMP device metrics.

Splunk Enterprise Security Integration

CIM-compliant output maps directly to Splunk ES correlation searches and notable event workflows. Enriched fields — user, threat score, GeoIP — are available immediately in ES risk-based alerting.

Splunk Observability Cloud (OTel)

NFO's OpenTelemetry output delivers enriched NetFlow and SNMP device metrics directly to Splunk Observability Cloud — bringing network-layer visibility into the same platform as your application and infrastructure telemetry.

Before NFO vs. After NFO in Splunk

Scenario	Without NFO	With NFO
Splunk daily ingest volume	High — raw flows, full volume	**80–90% lower — deduplicated & aggregated**
Data context in Splunk	IP addresses only	**User, threat, GeoIP, VM — enriched at ingest**
CIM compliance	Manual field mapping required	**Automatic — NFO normalizes inside the pipeline**
ES correlation searches	Partial — missing network context fields	**Enriched with user identity, threat scores, and GeoIP — network fields carry full investigative context for ES correlation**
SNMP device health in Splunk	Separate tool, separate dashboard	**Same pipeline, same Splunk app**
Splunk search performance	Slow — high volume, low cardinality	**Fast — smaller volume, high-value events**

SYSTEM REQUIREMENTS	SIZING GUIDANCE
Linux: RHEL 7+, Rocky Linux 8+ **Windows:** Windows Server 2019, 2022, 2025 *Software-only — no proprietary hardware required.*	**Entry level:** 2 CPUs, 8 GB RAM, 20 GB disk **Throughput:** 300K+ flows/sec at entry-level sizing **Scale-out:** Add instances; NFO Central manages distributed deployments with no throughput ceiling

See how much NFO can cut your Splunk bill.

Start a free trial with your own network or schedule a technical demo with a NetFlow Logic engineer.

Start Free Trial
netflowlogic.com/free-trial

Schedule a Demo
netflowlogic.com/request-a-demo