SUPERNETFLOW – REINVENTING NETWORK OBSERVABILITY WITH EBPF

OVERVIEW

Canopus Networks is a cloud native network observability company specializing in extracting deep, real-time insights from high-speed networks. To address the growing demands of telecommunications and enterprise operators for visibility into user experience and application behavior, Canopus developed "SuperNetflow," a powerful observability pipeline that leverages eBPF to monitor and classify traffic at scale. Using eBPF's programmable in-kernel processing capabilities, Canopus delivers detailed insights on network flows – including application identification and quality-of-experience (QoE) metrics – across carrier, fixed-line, and enterprise environments.

CHALLENGE

Network operators face increasing pressure to improve service quality and monetize infrastructure while coping with rising encryption, massive data volumes, and architectural diversity. Traditional DPI appliances and user-space packet processors like DPDK imposed trade-offs: costly hardware, limited cloud native compatibility, and high compute demands. Canopus needed a portable, programmable, and cost-efficient method to extract application layer intelligence, including video resolution, gaming lags, and conferencing glitches, in real time, without sacrificing performance or scalability.

Operators were also struggling with:

- Customer churn and complaints due to poor app experiences (e.g., WhatsApp, YouTube).
- Lack of visibility into streaming and gaming behavior to guide business decisions.
- Inability to track application experience post-network changes.
- Incompatibility of existing tools with cloud native and 5G SA/NFV architectures.

SOLUTION

Canopus chose eBPF for its unique ability to extract rich network data in-kernel with low compute overhead, high throughput, and cloud native deployment flexibility. They implemented a stateful, programmable packet processing pipeline using eBPF's XDP (eXpress Data Path), enabling on-the-fly traffic classification and telemetry extraction.

The core of their solution, SuperNetflow, uses eBPF maps and flow-aware telemetry to:

- Embed sensors in bare metal, VMs, or containers.
- Dynamically suppress or truncate traffic based on classification (e.g., Netflix flows).
- Extract behavioral time-series metrics (e.g., byte patterns, DNS, QoE scores).
- Minimize compute load by reducing unnecessary user-space handoff.
- Leverage shared flow maps between kernel and user-space for application logic.

RESULTS

Canopus's eBPF-based observability pipeline delivered transformative benefits across several dimensions:

- **Portability**: Sensors can be deployed anywhere, on-prem, cloud, or at the edge, without requiring hardware appliances. This enabled widespread adoption across mobile, fixed-line, and enterprise networks.
- **Programmability**: eBPF allowed Canopus to adapt telemetry collection on-demand, minimizing data export volumes while maximizing insight. With the final data stream out of the eBPF sensors only **0.1% of the previous bitrate**.
- **Scalability**: Achieved **1 Gbps per core processing**, with full-stack L2–L7 visibility. In production, Canopus systems reached **up to 400 Gbps traffic inspection**.
- **Business Impact**:
  - **Reduced server footprint 3x** compared to DPI appliances.
  - Enabled telcos to detect app-level performance degradations across the network at the subscriber level.
  - Helped customers build a data moat for business operations and strategy.

WHY eBPF?

eBPF provided Canopus with the flexibility, performance, and ecosystem support necessary to build a future-ready observability platform:

- Extract high-fidelity telemetry directly in the kernel.
- Avoid dependencies on third-party libraries and specialized NICs.
- Eliminate hardware lock-in, enabling pure software-based deployment.
- Align with Kubernetes and cloud native operational models.
- Increase telemetry scalability and performance of high throughput networks.

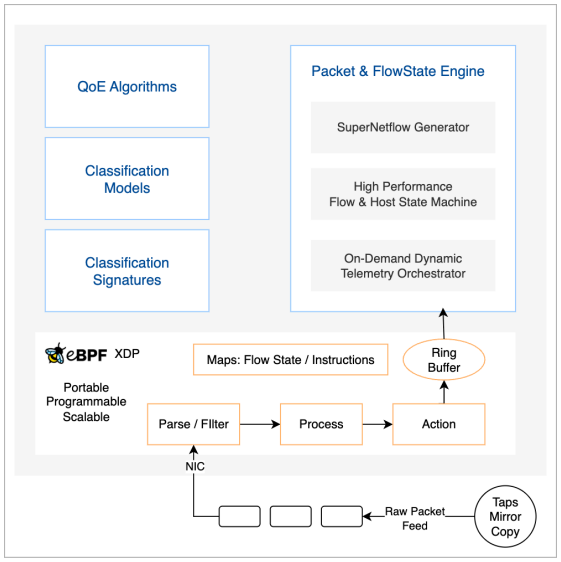

Fig 1. SuperNetflow Architecture

> "eBPF gave us the ability to process network traffic with the same principles as P4 programmable switches – but entirely in software, with better flexibility, cost-efficiency, and scalability."
>
> **-Himal Kumar, CTO, Canopus**