DYNANIC Silicom Denmark A/S FPGA SOLUTIONS

Experience Unmatched Network Efficiency and FPGA Excelence

Accelerating AI-based Network Monitoring

Problem statement

As network traffic grows in volume and complexity, modern security and operational monitoring systems require detailed insights at line rate. Traditional approaches struggle to **keep up with the real-time extraction of advanced features from network traffic**. Extracting relevant metadata—such as packet lengths, interpacket gaps, entropy, or identifying complex communication patterns—places a **significant burden on CPUs and servers**.

To manage this load, **many systems rely on packet sampling**, which reduces processing overhead but significantly limits network visibility. **Sampling introduces blind spots in traffic analysis**, making it difficult to detect subtle anomalies, low-rate threats, or new types of suspicious behavior, especially in encrypted flows.

Without hardware support, **AI-based traffic analysis is constrained by performance bottlenecks**, excessive resource usage, and limited scalability. This restricts the ability to **analyze traffic effectively**, especially at **100G or 400G link speeds**.

DYNANIC addresses this challenge by accelerating feature extraction and streamlining traffic pre-processing in hardware, enabling scalable and cost-efficient AI-based monitoring.

Solution description

DYNANIC offers a programmable FPGA-based SmartNIC designed to efficiently preprocess, analyze, and classify high-speed network traffic for AI-powered monitoring systems. The **SmartNICs such as Silicom ThunderFjord (FB2CDG1@AGM series), N6010, or FB2CG(HH)@KU15P SmartNICs operates directly on live traffic**, ensuring the **wire-speed throughput of 100/200/400G links** is maintained while reducing the amount of traffic forwarded to host CPUs or GPUs.

For more information visit
www.dyna-nic.com and www.silicom.dk

Unlike traditional monitoring solutions, **DYNANIC avoids packet sampling entirely.** This ensures **complete visibility across all network traffic**, enabling more accurate analysis and more accurate, AI-driven decision-making.

The DYNANIC AI-based monitoring function relies on a carefully structured two-stage processing architecture that ensures maximum efficiency, scalability, and readiness for integration with AI-powered monitoring systems. This architectural model is **designed to handle extremely high volumes of network traffic** without compromising speed or insight.

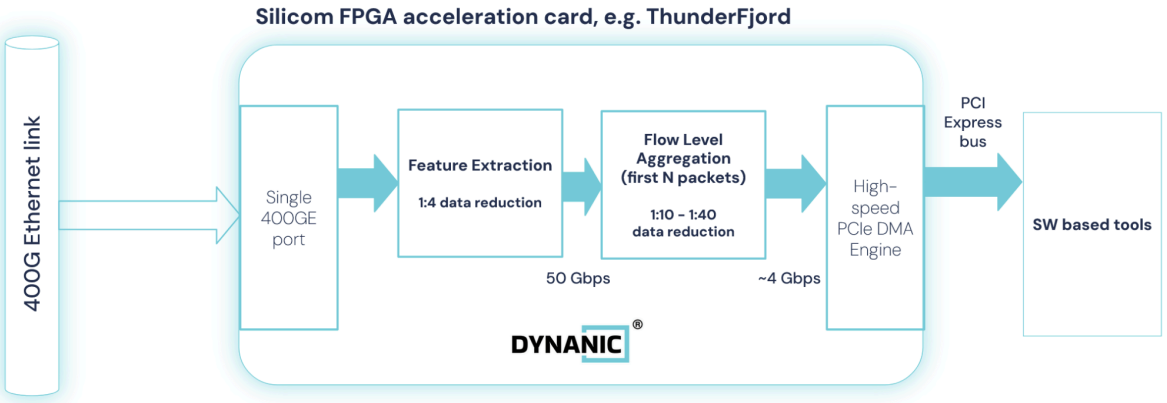

The **first stage focuses on per-packet feature extraction**. As packets enter, dedicated FPGA logic inspects them in real-time, parsing key fields such as **IP addresses, ports, protocol types, packet lengths, TCP flags, timestamps, and selected payload bytes**. Lightweight metrics such as **entropy** are also computed to help characterize the nature of the traffic. This stage reduces the data rate significantly (approx. 1:4) by replacing raw packets with **compact metadata representations** that retain all essential information for downstream analysis.

In the **second stage**, this **packet-level metadata is aggregated into flow-level records** through a hardware-based connection tracking table. This stateful mechanism correlates packets belonging to the same network flow and computes **aggregated metrics such as byte/packet counts, flow duration, packet size distributions, inter-arrival gaps, TCP flag activity, and directional symmetry**. These flow summaries are much smaller and easier to interpret than raw packet streams, allowing the host system to focus entirely on AI inference and behavioral analysis.

By combining both stages, DYNANIC achieves an **end-to-end optimization that dramatically reduces the volume of data** sent to the host system. Instead of forwarding millions of raw packets per second, the **solution produces a compact stream of enriched and aggregated flow records**. This output is precisely tailored to the needs of modern AI models, which rely on structured metadata rather than raw traffic to detect patterns, classify applications, and identify anomalies.

This architectural approach allows the **CPU to focus exclusively on high-value AI tasks** such as inference, threat classification, and model retraining. The system no longer wastes cycles on packet parsing or statistics computation, which are offloaded to the FPGA card. By leveraging **Silicom's high-performance FPGA SmartNICs — such as the ThunderFjord (FB2CDG1@AGM series) or N6010,** the solution benefits from advanced hardware acceleration capabilities. These cards feature cutting-edge FPGAs with high-speed interfaces and substantial on-board memory (e.g., DDR4 or 32GB HBM2e), enabling efficient offloading of compute-intensive tasks. The reconfigurable nature of these FPGAs allows the feature extraction pipeline to be easily adapted to support evolving monitoring requirements and AI models, **making the solution highly flexible and future-proof.** As a result, even at extremely high traffic volumes, such as 400 Gbps or beyond, DYNANIC enables scalable and cost-effective AI-driven monitoring without requiring extensive server infrastructure.

Key Advantages

- **High-Speed Feature Extraction at Line Rate**: DYNANIC enables AI-based monitoring even on encrypted traffic by extracting meaningful features and statistical patterns even at 400G networks.
- **Customizable and Evolvable Feature Sets**: FPGA-based implementation on Silicom SmartNICs allows continuous expansion of feature sets to match evolving AI model requirements, offering long-term adaptability and investment protection.
- **Massive CPU/GPU Offload**: Reduces the amount of traffic forwarded to CPU or GPU systems by 40x–160x, minimizing processing latency, energy consumption, and system costs.
- **Inline Intelligence**: Real-time feature detection and filtering on the NIC reduces time-to-insight for security and performance monitoring, especially in sensitive environments.
- **Self-Updating Dataset Collection**: A built-in sampling engine enables ongoing training data collection based on custom rules, helping to keep AI models current and preventing data drift.

Value for Customers

Telco operators benefit from the ability to **monitor encrypted or obfuscated traffic** at the edge without payload decryption, **improving visibility** while complying with privacy regulations.
Data Centers and Cloud Providers can **reduce the number of monitoring servers required by 50–70%**, saving on space, power, and cooling while boosting analytics throughput.

AI brings a new level of visibility to network traffic:

- Application protocol fingerprinting (even in encrypted streams)
- Communication-type classification
- Threat detection (e.g., covert channels, beaconing patterns)
- Enrichment of SIEM and XDR systems with high-value metadata

A server equipped with one or more Silicom FPGA-based SmartNICs running DYNANIC enables real-time, AI-powered monitoring with massive performance and efficiency gains. The solution bridges the gap between hardware acceleration and evolving AI algorithms, ensuring scalability, flexibility, and cost-effective insights for next-generation network observability.

Feel free to contact us at:

info@dyna-nic.com or contactus@silicom.dk