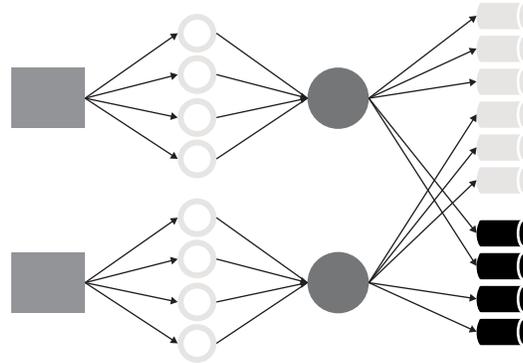


High-Performance Packet Path Accelerator

Reservoir Labs | Technology Brief



mCore

As the capacity of network links and multicore CPU/GPU processors continue to increase, a new bottleneck is emerging at the network interface that prevents performing network security at scale. Today a single fiber link can carry about 100 terabits per second of data traffic, while state-of-the-art multicore CPU/GPU processors can aggregately scale at a similar rate. Yet, traditional network stacks implemented by the operating system are limited to rates up to 100s of gigabits per second. Further, as computational parallelism has increased, this bottleneck has dramatically worsened.

This situation has been recognized by both the scientific community and industry. For instance, state-of-the-art artificial intelligent applications using massive amounts of data are finding the network interfaces to be a major obstacle to achieving scale. Similarly, companies that manage very large-scale distributed systems require parsing packets at the rate of multiple terabits per second in order to secure and identify potential cyber attacks.

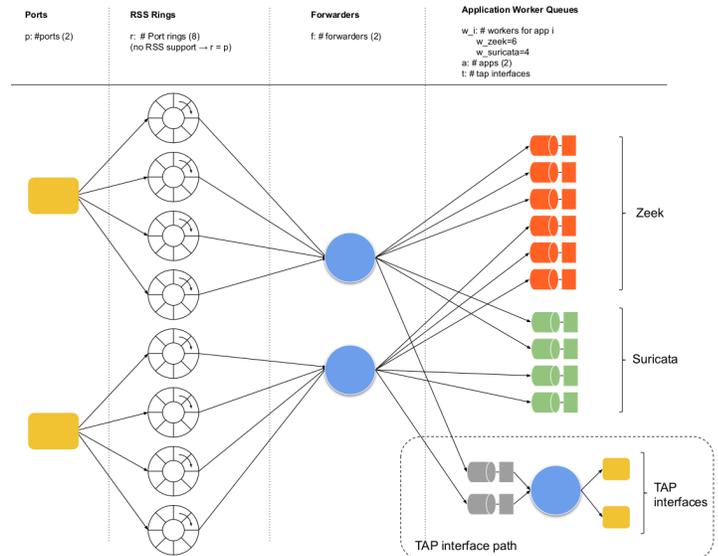
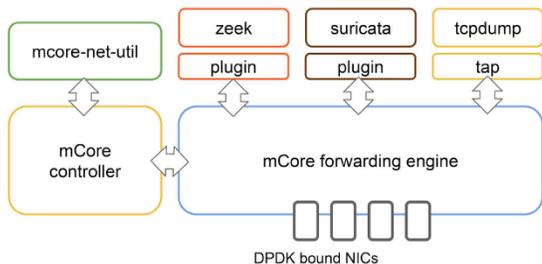
To address this gap, Reservoir Labs has designed and patented mCore, a suite of cybersecurity-aware data structures and algorithms designed to address and resolve the network interface bottleneck.

Experimental benchmarks demonstrate that at input rates of 10Gbps, aggregately these optimizations increase application performance up to 500% while packet drops are reduced up to 200%.

Features

- High-Performance Data Structures
- Reduce Memory Contention
- Minimize Back-Pressure
- Maximize Security Metadata
- New Packet Polling Algorithm
- High Speed Software APIs
- Leverage State-of-the-Art NICs
- Support for DPDK High-Performance Packet Path

mCore in Action



—○ mCore Architecture

In the mCore architecture, packets arriving from DPDK-bound network ports are injected into the mCore forwarding engine, which is responsible for delivering them to the application layer enabling a high-performance packet path. The forwarding engine provides the following key high-performance features:

- Kernel Bypass
- Zero Packet Copy
- Lockless Data Structures
- NUMA Affinity
- CPU/core Affinity/Pinning
- Multi-Core Elastic Scalability

—○ Packet Forwarding Engine

The architecture is designed to enable these features:

- Scalability to increase the number of application workers while avoiding bottlenecks across the pipeline
- Flexibility to assign a different number of worker processes to different applications based on their computational requirements
- End-to-end lockless architecture to ensure that no shared data structure needs to be locked, which would impose a performance penalty due to memory contention
- Ensure NUMA coherency to avoid performance penalties due to memory accesses on remote NUMA nodes
- End-to-end zero packet copies to reduce memory and CPU bottlenecks
- Runs on DPDK to leverage hardware-aided packet acceleration