




PRODUCT BRIEF

Link-Inline™ Software for Napatech FPGA-based SmartNICs



NT200A02 SmartNIC

5G Telecom User Plane Traffic Challenges

As Communications Service Providers and enterprises scale up the deployment of their 5G networks, it becomes critical to optimize the Return on Investment for their infrastructure. With 5G User Plane Function (UPF) representing a significant part of the overall workload for the 5G Packet Core, it's important to leverage solutions that maximize the number of users supported per server and thereby minimize the overall cost per user.

By offloading the UPF data plane processing from the server CPU to an integrated hardware/software SmartNIC solution, Napatech delivers industry-leading performance for the UPF function. This offload solution is based on industry-standard APIs and cloud-native orchestration, enabling it to be easily integrated into complete 5G Core products offered by telecom software vendors.

Cyber Security Processing Challenges

Network security architects are seeing requirements for their solutions quickly changing with the explosion in network throughputs while at the same time, the threat landscape is continuously evolving and becoming more complex and sophisticated. Stateless security solutions are no longer adequate to identify and block threats. Inline networking and security solutions require complete and stateful awareness of all users and applications at throughputs topping 100Gbps.

To support these requirements, network infrastructures need to contain more intelligence with deeper inspection of traffic at increasing line rates. With this need for inline stateful flow processing, application awareness, content

inspection, and security processing, the amount of compute power to meet these increasing line rates grows exponentially.

Stateful Flow Management and Offload

To maintain performance at high speeds and address all of these challenging issues, Link-Inline™ software offloads packet and flow-based processing to reconfigurable FPGA-based SmartNICs. The SmartNIC performs flow classification and identification on ingress and maintains state for each packet of a flow. For known flows, action processing is dynamically handled completely in the SmartNIC and all other packets are forwarded to the application for additional analysis to minimize the load on user-space applications. Link-Inline™ Software additionally provides the ability to dynamically identify and direct data flows into specific CPU cores based on the type of traffic being analyzed. Link-Inline™ software is tightly coupled to x86 cores for inline network and security applications. Per-flow match/action processing in HW gives control back to the user providing additional computation to the application by reducing the amount of data needed for processing as certain flows or protocols that no longer need monitoring and can be blocked or forwarded in hardware.

Napatech's Link-Inline™ software accelerates standard Linux applications and provides open APIs for development and integration of inline network applications. The solution significantly reduces host CPU utilization and solution latency by offloading complex flow classification and packet processing to the SmartNIC.

Stateful flow management

- Up to 140 million flows
- Learning rate: > 1.5 million flows/sec
- Flow match/actions:
 - Drop
 - Fast forward (inline / hairpin)
 - RSS (CPU load distribution)
 - GTP encapsulation/decapsulation (1k L2, L3 or L4 setups)
 - Tunnel ID configured per flow
 - MBR (Maximum Bit Rate) policing
 - DSCP (Differentiated Services Code Point) tagging
 - NAT (Network Address Translation)
 - Flow mirroring
- Flow termination: TCP protocol, timeout, application-requested
- Flow records: Rx packet/byte counters and TCP flags, delivered to application at flow termination
- Configurable flow definitions based on 2-, 3-, 4-, 5- or 6-tuple

Pre-filtering

- Configurable 2-, 3-, 4-, 5- or 6-tuple, enabling up to 36,000 IPv4 or up to 8,000 IPv6 addresses
- 864 (32-bit) wildcard entries
- General purpose filters: Pattern match, network port, protocol

CPU load distribution

- Load distribution based on hash key, filter or per flow
- Hash keys, calculated on 5-tuple from inner or outer headers

Rx Packet Processing

- 128 Rx queues
- Multi-port packet merge, sequenced in time stamp order
- L2, L3 and L4 protocol classification
 - L2: Ether II, IEEE 802.3 LLC, IEEE 802.3/802.2 SNAP
 - L2: PPPoE Discovery, PPPoE Session, Raw Novell
 - L2: ISL, 3x VLAN, 7x MPLS
 - L3: IPv4, IPv6
 - L4: TCP, UDP, ICMP, SCTP
- Tunneling support
 - GTP, IP-in-IP, GRE, NVGRE, VxLAN, Pseudowire, Fabric Path

Tx Packet Processing

- 128 Tx queues

Advanced Statistics

- Extended RMON1 per port
- Packets and bytes per flow
- Flow programming statistics

Monitoring Sensors

- PCB temperature level with alarm
- FPGA temperature level with alarm and automatic shutdown
- Temperature of critical components
- Individual optical port temperature and light level with alarm
- Voltage or current overrange with alarm
- Cooling fan speed with alarm

Supported OS and Orchestration

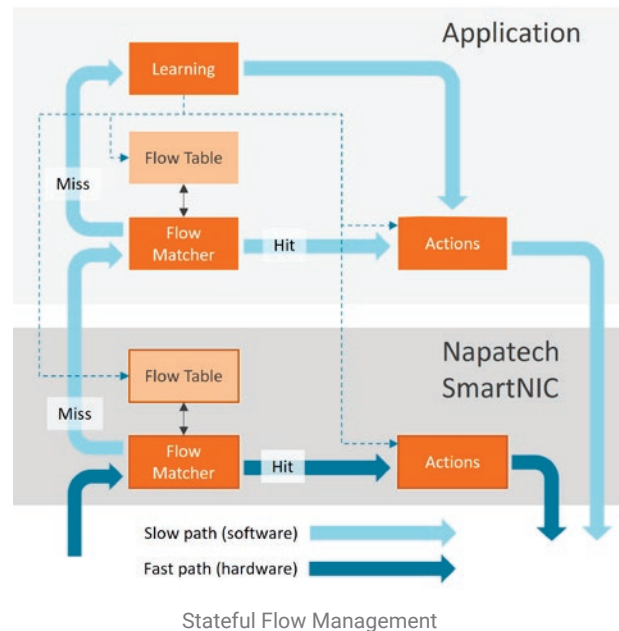
- Linux kernel 5.17 (64-bit)
- Kubernetes

Supported APIs

- DPDK v. 21.11.1, RTE_FLOW, RTE_METER

Supported Hardware and Transceivers

- NT200A02:
 - 100GBASE-SR4, SR-BiDi, LR4



NAPATECH RECONFIGURABLE COMPUTING

NAPATECH.COM

Disclaimer: This document is intended for informational purposes only. Any information herein is believed to be reliable. However, Napatech assumes no responsibility for the accuracy of the information. Napatech reserves the right to change the document and the products described without notice. Napatech and the authors disclaim any and all liabilities. Napatech is a trademark used under license by Napatech A/S. All other logos, trademarks and service marks are the property of the respective third parties. Copyright © Napatech A/S 2023. All rights reserved.