




PRODUCT BRIEF

Link™ Inline Software for Napatech FPGA-based SmartNICs



Link™ NT200A02 SmartNIC

Network and 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 via open APIs 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.

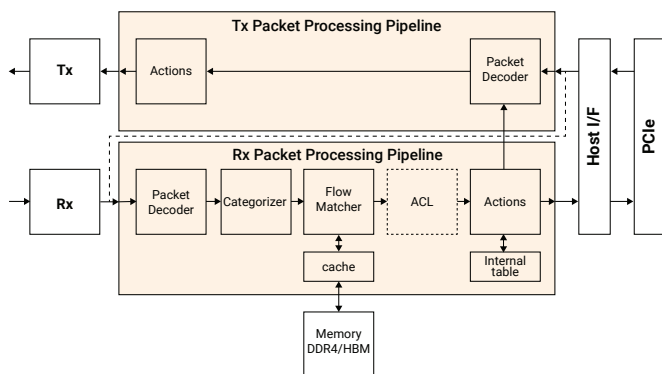
Napatech's Link™ Inline software accelerates standard Linux applications and provides open APIs for development and integration of inline network and security applications. The solution significantly reduces host CPU utilization and solution latency by offloading complex flow classification and packet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

Flow Classification

Flow classification is the ability to decode a packet where the parameters used to identify a flow are captured based on known IP parameters. If the flow is known, the action assigned to the flow can be based on policy or dynamically assigned via the application parameters including the ability to perform one or more actions on the packet. If the flow is unknown, it is forwarded to the application for further processing. Rules and open APIs provide the application the ability to dynamically change the action once a flow has been identified, application known, and threat status verified.

Rx Packet Processing

- Line rate Rx up to 100 Gbps for packet size 64 – 10,000 bytes
- 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
- L2 and L3/L4 (IP/TCP/UDP) Tx checksum generation
- L2 and L3/L4 (IP/TCP/UDP) Rx checksum verification
- Tunneling support
 - GTP, IP-in-IP, GRE, NVGRE, VxLAN, Pseudowire, Fabric Path
- General purpose filters
 - Pattern match, network port, protocol, length check, error conditions
- Stateful flow management (NT100A01 and NT200A02)
 - Configurable flow definitions based on 2-, 3-, 4-, 5- or 6-tuple
 - Up to 160 million bidirectional IPv4 or IPv6 flows (NT200A02)
 - Up to 100 million bidirectional IPv4 or IPv6 flows (NT100A01)
 - Learning rate: 3 million flows/sec
 - Flow match/actions: forward to specific host Rx queue, drop, fast forward to network port, update metrics in flow record
 - Flow termination: TCP protocol, timeout, application-requested
 - Flow records: Rx packet/byte counters and TCP flags, delivered to application at flow termination
- Hash keys
 - Custom 2 × 128 bits and 2 × 32 bits with separate bit masks
 - Symmetric hash keys
 - Protocol field from inner or outer headers
- Load distribution
 - Hash key, filter-based or per flow
 - To local CPU cores via host buffers/queues
- Packet descriptors and metadata
 - PCAP and Napatech descriptor formats
 - Time stamp, network port ID, header offsets
 - Hash key, color/tag
 - 64-bit pointer for flow lookup
 - 64-bit correlation key with maskable fields (packet fingerprint)
 - Protocol and error information
- Slicing
 - Dynamic offset or fixed offset from start or end of packet



Tx Packet Processing

- Line rate Tx up to 100 Gbps for packet size 64 – 10,000 bytes
- Per-port traffic shaping
- Port-to-any-port forwarding

Latency

- Port to port: < 5 us
- Port to host: < 30 us

Host Buffers and Queues

- Rx queues: 128, Tx queues: 128
- Rx buffer size: 16 MB – 1 TB, Tx buffer size: 4 MB

Advanced Statistics

- Extended RMON1 per port
- Packets and bytes per filter/color and per stream/queue

Time Stamping and Synchronization

- OS time, PPS and IEEE 1588-2008 PTP V2
- NT-TS synchronization between Napatech SmartNICs
- Time stamp formats: Unix 10 ns, Unix 1 ns, PCAP 1 us, PCAP 1 ns
- Rx time stamp with 1 ns resolution

Monitoring Sensors

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

Supported OS

- Linux kernel 3.0 – 3.19, 4.0 – 4.18, 5.0 – 5.11 (64-bit)

Supported APIs

- libpcap v. 1.7.3, 1.8.1, 1.9.0
- DPDK v. 20.11
- NTAPI (Napatech API)

Supported Hardware and Transceivers

- NT200A02:
 - 100GBASE-T, SX, LX, ZX
 - 10GBASE-SR, CR, LR, ER and breakout to 4x10GBASE-SR, CR, LR
 - 25GBASE-SR, LR, LR-BiDi and breakout to 4x25GBASE-SR, LR
 - 40GBASE-SR4, SR-BiDi, CR4, LR4
 - 100GBASE-SR4, SR-BiDi, LR4
- NT100A01:
 - 100GBASE-T, SX, LX, ZX
 - 10GBASE-SR, CR, LR, ER
 - 25GBASE-SR, LR, LR-BiDi

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 2021. All rights reserved.