napa:tech;
RECONFIGURABLE COMPUTING

PRODUCT BRIEF

Link-Capture™ Software for Napatech FPGA-based SmartNICs

The Necessity of Full Packet Capture
Today, a massive selection of tools is available to help engineers and administrators manage and secure their networks. Still, few capabilities are as fundamental to this task as packet capture.

Full packet capture provides an accurate, real-time view of what is happening within a network infrastructure. It provides organizations with the ability to recreate network events with high fidelity for verification and validation of architectural changes, measure network performance, troubleshoot issues, and perform forensic analysis to determine the impact of network breaches. However, as network speeds continue to increase, existing capture solutions are struggling to keep up.

Link-Capture™ Software
In helping enterprises, network operators and equipment vendors overcome this challenge, Napatech has developed Link-Capture™ Software. This technology can immediately improve an organization's ability to monitor and react to events that occur within its network infrastructure. The speed and accuracy of the software's programmable logic enables:

- **Line-Rate Capture and Replay**
  Link-Capture™ Software is ideal for performing high-speed packet capture with nanosecond time-stamping and replay with precise inter-frame gap control, which is critical when replaying captured traffic for troubleshooting or simulation of traffic flows.

- **Stateful Flow Management**
  To maintain capture and analysis performance at high speeds, it is important to identify and direct traffic flows immediately upon ingress to minimize the load on user-space applications. Link-Capture™ Software provides the ability to dynamically identify and direct data flows into specific CPU cores based on the type of traffic being analyzed. Per flow match/action 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 no longer need monitoring and can be blocked in hardware.

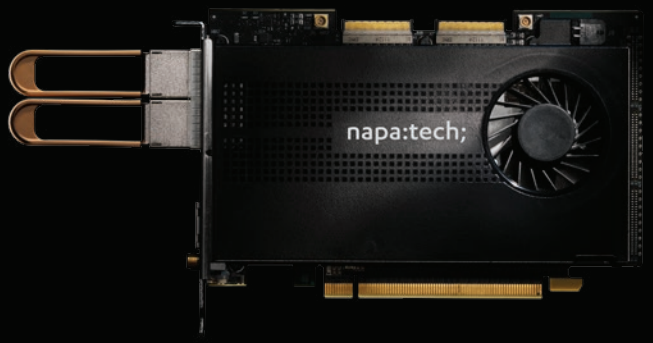

Napatech Link-Capture™ Software supports a broad range of applications and use cases. Where standard Network Interface Cards (NICs) suffer from intolerable packet loss for the target applications, Napatech guarantees line rate throughput with zero packet loss and replay for all packet sizes.

Multiplied Application Performance
Link-Capture™ Software has been benchmarked across a wide range of third-party, commercial and open source networking and cybersecurity applications. Common to these is the unconditional requirement for line rate throughput for all packet sizes, with 100% lossless packet forwarding and capture, for a multitude of sessions, users and flows.

With Link-Capture™ Software, the performance improvements are outstanding, delivering more than triple the performance over servers with standard NIC configurations. This means a third of the required server resources to run the same application.

Key Features
- Zero packet loss under all conditions
- Full throughput up to 100 Gbps bi-directional
- Nanosecond time-stamping and packet merge
- Up to 140 million flows with stateful match/action
- Flow records with metrics for both directions
- PCAP and DPDK API support

FEATURES AND APPLICATIONS

Rx Packet Processing
- Line rate Rx up to 100 Gbps for packet size 64 – 10,000 bytes
- Zero packet loss
- 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
- Flow filtering
  - Configurable flow definitions based on 2-, 3-, 4-, 5- or 6-tuple
  - Up to 36,000 IPv4 or up to 7,500 IPv6 2-tuple flows
  - Flow match/actions: forward to specific host Rx queue, drop, fast forward to network port, select packet descriptor type, slice, strip headers, mask
- Stateful flow management
  - Configurable flow definitions based on 2-, 3-, 4-, 5- or 6-tuple
  - Up to 140 million bidirectional IPv4 or IPv6 flows
  - Up to 90 million bidirectional IPv4 or IPv6 flows
  - 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
  - Remotely via VLAN tagging over egress port
  - CPU Socket Load Balancing
- 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
- IP fragment handling
  - First level IP fragmentation
  - Filter actions on inner header fields applied to all fragments
- Deduplication
  - Configurable action per port group: discard or pass duplicates
  - Duplicate counters per port group
- Slicing
  - Dynamic offset or fixed offset from start or end of packet
- Header stripping
  - Removal of protocol layers between outer L2 and inner L3 headers
- Packet masking
  - Zeroing of 1 – 64 bytes at dynamic or fixed offset

Note: Feature set varies depending on the SmartNIC model.

Tx Packet Processing
- Line rate Tx up to 100 Gbps for packet size 64 – 10,000 bytes
- Replay as captured with nanoseconds precision
- Per-port traffic shaping
- Port-to-any-port forwarding

Rx Buffer Capacity
- NT400D11, NT200A02: 12GB
- NT100A01: 8GB
- NT50B01: 10GB
- NT40A11: 4GB

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
- 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
- Tx time stamp inject
- 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.x, 4.x and 5.x (64-bit)
- Windows server 2022 (64-bit) and Windows 11 (64-bit)

Supported APIs
- libpcap v. 1.8.1, 1.9.1, 1.10.0 and WinPcap v. 4.1.3
- DPDK v. 21.11 LTS
- NTAPI (Napatech API)

Supported SmartNICs and Transceivers
- NT400D11:
  - 100GBASE-SR4, SR-BiDi, LR4, PSM4
- NT200A02:
  - 1000BASE-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 and NT50B01:
  - 1000BASE-T, SX, LX, ZX
  - 10GBASE-SR, CR, LR, ER
  - 25GBASE-SR, LR, LR-BiDi
- NT40A11:
  - 1000BASE-T, SX, LX, ZX
  - 10GBASE-T, SR, CR, LR, ER

DN-1202 Rev. 14