

Accelerators for Smarter Data Delivery

N-ANL9 FEATURE OVERVIEW

PLUG-AND-PLAY INTELLIGENT ACCELERATION

The Napatech accelerator family supports a common feature set and driver software architecture allowing plug-and-play support for any accelerator combination.

The intelligent feature set off-loads processing and analysis of Ethernet data from application software while ensuring optimal use of the standard server's resources leading to effective application acceleration.

FEATURE DESCRIPTIONS

Full Line-Rate Packet Capture

- Zero-loss packet capture for all frame sizes

Optimized Packet Buffers

- Low CPU load on the host server

Onboard Packet Buffering

- Ensures packet delivery during congestion

HW Time Stamping

- 1 ns time stamp resolution

Time Synchronization

- IEEE 1588-2008 PTP V2
- PPS
- OS time

Frame Processing

- Multi-port and multi-accelerator data merge
- Frame classification
- Frame and protocol information
- Fixed, dynamic and conditional slicing
- Ethernet FCS verification

Tunneling Support

- GTP, IP-in-IP, GRE, NVGRE and VxLAN tunneling support
- Inner tunnel slicing and filtering of packet headers and payloads

IP Fragment Handling

- Identification and distribution of IP fragments

Flow Identification

- Flow identification based on hash keys
- Multiple protocol-specific hash keys
- Dynamic hash key selection

Frame and Flow Filtering

- Configurable filters
- Port and protocol filters
- IP group match filters
- Pattern compare filters
- Coloring

Intelligent Multi-CPU Distribution

- Configurable data distribution
- QPI bypass for 2 x 100G performance solution

Advanced Statistics

- RMON1 port statistics
- Advanced statistics

Monitoring Sensors

- Temperature and power sensors

Host-based Transmission

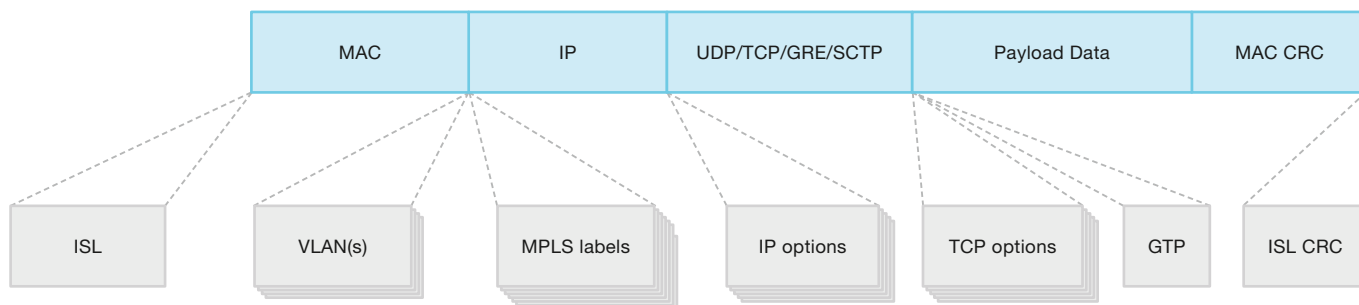
- 128 Tx Streams

In-Line Application Support

- Supports multi-core processing with up to 128 Rx/Tx streams per accelerator
- Zero copy transfer from Rx to Tx
- Single bit flip to select discard or forward packet

Flow Match

- Packet discard, forwarding of packet to a specific stream, packet coloring and slicing of packet controlled per flow by the application



The Napatech accelerators decode all frames regardless of encapsulations and can identify flows based on user defined header information

FULL LINE-RATE PACKET CAPTURE

Napatech accelerators are highly optimized to capture network traffic at full line rate, with almost no CPU load on the host server, for all frame sizes. Zero-loss packet capture is critical for applications that need to analyze all the network traffic.

If anything needs to be discarded, it is a matter of choice by the application, not a limitation of the accelerator.

OPTIMIZED PACKET BUFFERS

Napatech accelerators use a buffering strategy that allocates a number of large memory buffers where as many packets as possible are placed back-to-back in each buffer. Using this implementation, only the first access to a packet in the buffer is affected by the access time to external memory. Due to cache pre-fetch, the subsequent packets are already in the level 1 cache before the CPU needs them. As hundreds or even thousands of packets can be placed in a buffer, a very high CPU cache performance can be achieved leading to application acceleration.

The flexible server buffer structure supported by Napatech accelerators can be optimized for different application requirements:

- Applications needing short latency can have frames delivered in small chunks, optionally with a fixed maximum latency.
- Applications without latency requirements can have data delivered in large chunks, providing more effective server CPU processing of the data.

The customer can configure the size of the host buffers from 16 MB to 1 TB.

ONBOARD PACKET BUFFERING

Onboard accelerator memory buffers allow packets to be buffered during micro-burst or PCI Express bus congestion situations.

TIME STAMP

High-precision time stamps, sampled with 1 nanosecond resolution, are applied to all frames received by the accelerator. The accelerator also supports insertion of a high-precision 64-bit time stamp in frames being transmitted. The time stamp is inserted at a user-defined offset. Support is provided for 3 different 64-bit time-stamping formats:

- Native UNIX format with 10 ns resolution
- 2 PCAP formats with 1 ns or 1000 ns resolution

TIME SYNCHRONIZATION

The accelerator time-stamping can be synchronized to that of another accelerator or to external sources:

- Synchronization with GPS antennas and other external PPS time sources
- Synchronization with PTP time masters
- Accelerator-to-accelerator HW time synchronization using daisy-chaining of accelerators
- OS time synchronization with dynamic drift adjustment

ONBOARD IEEE 1588-2008 (PTP V2) SUPPORT

Onboard IEEE 1588-2008 support enables connection to a PTP network and synchronization to a PTP master directly from a Napatech accelerator.

Napatech accelerators with onboard IEEE 1588-2008 support include Default, Telecom (ITU-1 G.8265.1) profile support. Accelerators can be used as slave in Telecom and default profiles. These accelerators also provide the possibility of specifying a Packet Delay Variation (PDV) filter to improve time synchronization in a PTP unaware network.

Accelerators with onboard IEEE 1588-2008 support also provide additional features for enhanced time

synchronization. These include the ability to synchronize the host OS time to the time retrieved from PTP as well as the ability to generate a PPS output for synchronization of 3rd party equipment that does not support PTP directly. Napatech provides a complete PTP solution including hardware and software PTP stack.

FRAME PROCESSING

MULTI-PORT AND MULTI-ACCELERATOR PACKET MERGE

For accelerators with multiple ports, it is possible to merge packets received on multiple ports into a single data stream for further processing. This is useful for analyzing both receive and transmit directions of a connection, where the packets from the receive and transmit directions are received on separate ports. With Napatech Software Suite it is possible to merge data streams from multiple accelerators into a single data stream for analysis.

FRAME CLASSIFICATION

The accelerator frame decoder ensures protocol recognition of all major layer 2 to 4 protocols. The frame classification information is provided together with each frame including offset information for encapsulated protocol header and payload data. This provides a significant processing offload and acceleration of analysis.

Frame classification information can be used for:

- Finding protocol headers and payload data at dynamic locations
- Advanced filtering
- Dynamic selection of hash keys (per frame)
- Dynamic slicing
- Intelligent Multi-CPU Distribution

FRAME AND PROTOCOL INFORMATION

For each Ethernet frame received, frame and protocol information is made available to the host application. This accelerates processing of each frame. Information available includes:

- High-precision time-stamp information
- Protocol information: e.g. IPv4, IPv6, UDP, TCP, GRE, SCTP, EtherIP and GTP
- Encapsulation information: ISL, VLAN and MPLS
- Hash key information: hash key value and type
- Offsets to start of L3, L4 and L5 payload
- Coloring/tagging: Tags defined by filters
- Checksum error flags: Ethernet FCS

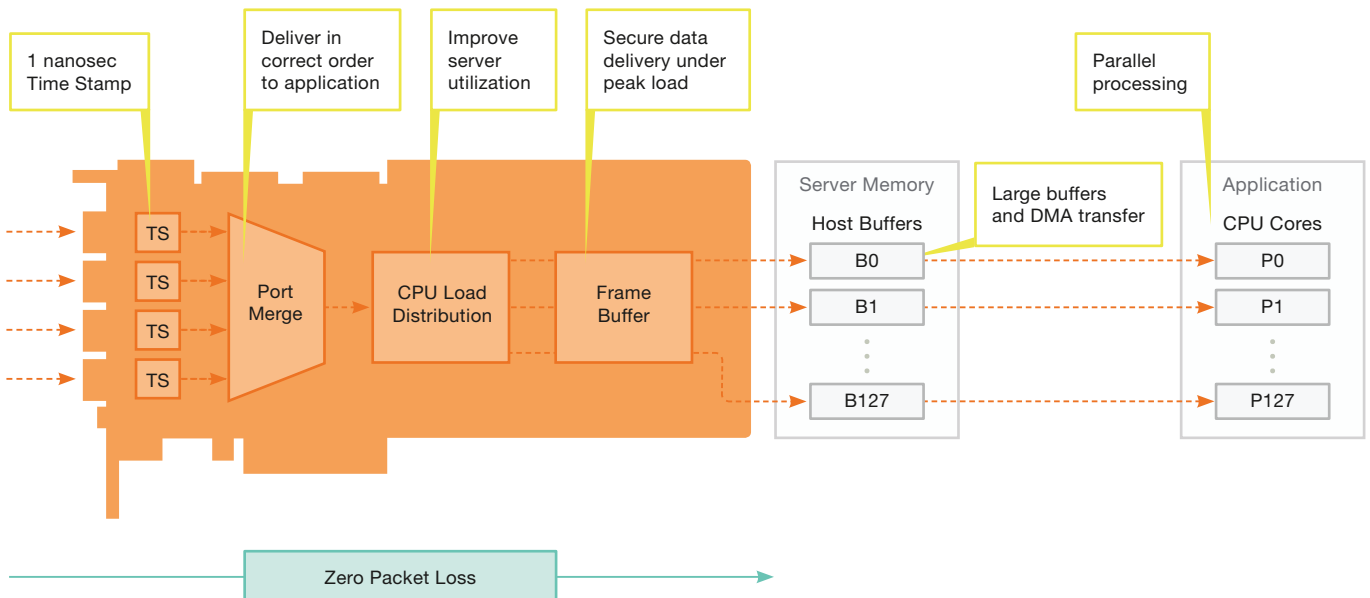
SLICING

With slicing it is possible to truncate frames so only essential information is analyzed. This reduces the processing load and thereby accelerates performance. Several types of slicing are supported:

- Fixed slicing: Truncation of frames to a maximum size or an offset from end of frame
- Dynamic slicing: Truncation of frames to a dynamic size relative to a specified protocol, e.g. IP payload + 16 bytes
- Conditional slicing: Fixed or dynamic slicing with properties based on frame decoding

ETHERNET FCS

For each received frame, the Ethernet Frame Check Sequence (FCS) is checked in hardware.



TUNNELING SUPPORT

GTP, IP-IN-IP, GRE, NVGRE AND VXLAN TUNNELING SUPPORT

Recognition of tunneling encapsulations enables effective CPU load distribution based on the contents of the tunnel rather than the tunnel itself. This is especially useful when analyzing traffic on telecom backbone networks using GTP or IP-in-IP, or on enterprise networks using IP-in-IP, GRE, NVGRE or VxLAN.

INNER TUNNEL FRAME PROCESSING

For tunneled traffic it is important to perform analysis on the content of tunnels rather than the tunnel itself including methods for reducing the amount of traffic to be analyzed through slicing and filtering:

- Slicing of encapsulated packet headers and payloads
- Filtering on encapsulated packet headers and payloads

IP FRAGMENT HANDLING

Napatech accelerators are capable of identifying fragmented IP packets and ensuring that associated fragments are distributed to the same host buffer for reassembly by the application. This helps accelerate the IP reassembly process for the application. Out of sequence fragments are also identified and handled by the accelerator.

Note: IP Fragment Handling does not apply to the NT200A01.

FLOW IDENTIFICATION

FLOW IDENTIFICATION BASED ON HASH KEYS

Frame classification information is used to calculate a hash key for each frame received. Frames with the same hash key can be treated as a flow and be processed accordingly.

MULTIPLE PROTOCOL-SPECIFIC HASH KEYS

Hash key calculation is based on combinations of specific header data. Napatech supports up to 17 different hash key types based on the type of protocols encapsulated in the Ethernet frame. Protocol-specific hash keys allows flows to be established based on the type of protocol to be analyzed.

DYNAMIC HASH KEY SELECTION

Since hash key types are protocol-specific, the accelerator can dynamically select the correct hash key type for calculation based on the information provided by the frame decoder.

FRAME AND FLOW FILTERING

CONFIGURABLE FILTERS

64 advanced programmable filters are available with an exceptional flexibility in the way they are specified, combined and controlled. They are configured by means of the easy-to-use Napatech Programming Language (NTPL). The

customer application can change the filters on the fly. The programmable filter logic is built on top of the advanced protocol decoding capabilities. This ensures that the application always gets the packets matching the requested protocol, even under very diverse conditions. Using a single filter, the accelerator can capture, for instance, all TCP/IP packets with a specific source IP address, even when these packets are ISL-, VLAN- or MPLS-encapsulated and/or contain IP/TCP options.

PORT AND PROTOCOL FILTERS

Filtering can be performed based on the port on which frames are received and the protocol information from the frame decoder.

IP MATCH FILTERS

Filtering can be performed based on the source and destination IP addresses of received packets including packets that are tunneled. Two pools of IP match addresses are available supporting both IPv4 and IPv6 protocols. For filtering of tunneled packets, one IP address pool can be used to match outer IP addresses while the other address pool can be used to match inner IP addresses. The IP addresses specified in the IP address pool can be individual IPv4 or IPv6 addresses or address ranges.

PATTERN COMPARE FILTERS

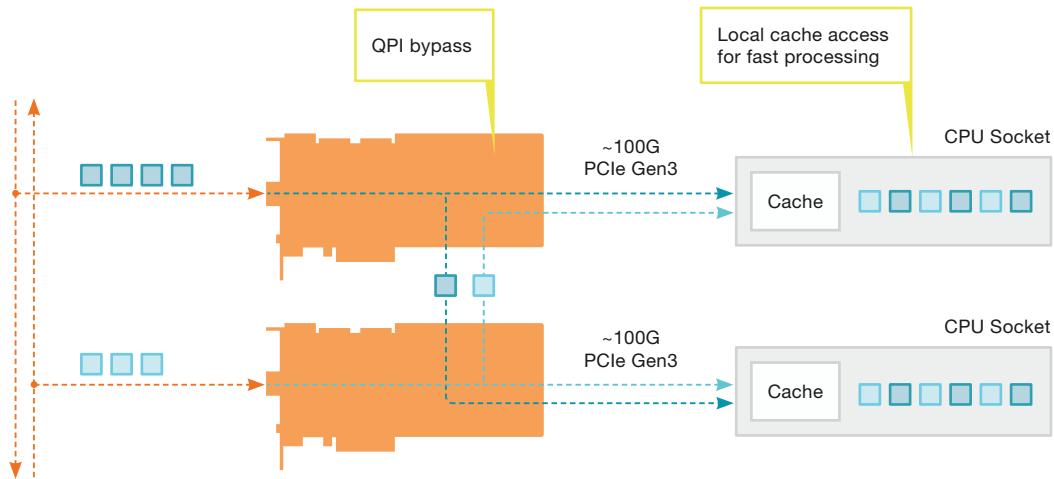
The pattern compare feature allows frames to be identified based on the application of a user defined data pattern with bit masks at up to 4 specific offsets in the frame. A single frame can be compared against multiple data patterns.

COLORING

Filtered frames can be tagged with a "color" ID identifying the filter that forwarded the frame. Color ID can be used to optimize applications performing different processing for different frame types.

FLOW MATCH FEATURES

- Packet discard, forwarding of packet to a specific stream, packet coloring and slicing of packet controlled per flow by the application
- Up to 18000 IPv4 or up to 2500 IPv6 2-tuple flows
- Up to 12000 IPv4 or up to 2500 IPv6 5-tuple flows
- Up to 36000 IPv4 or up to 7500 IPv6 source or destination IP address match entries
- Configurable flow definitions, based on 2, 3, 4 or 5-tuple
- Flow definition selected per frame category (e.g. IP protocol or TCP port number)
- Flow ID delivered to application in packet descriptor



2 x 100G Performance Solution

INTELLIGENT MULTI-CPU DISTRIBUTION

CONFIGURABLE DATA DISTRIBUTION

Multi-CPU buffer splitting enables accelerators to place captured frames in 1 to 128 host buffers. The customer can configure the size of the host buffers and how data is placed in the host buffers, based on results from the filter logic, port numbers and/or generated hash key values (flows).

The advanced multi-CPU buffer splitting functionality and the option for distributing traffic up to 128 CPU cores significantly improves the CPU cache performance, by always delivering the same flows and frame types to the same CPU cores.

QPI BYPASS FOR 2 x 100G PERFORMANCE SOLUTION

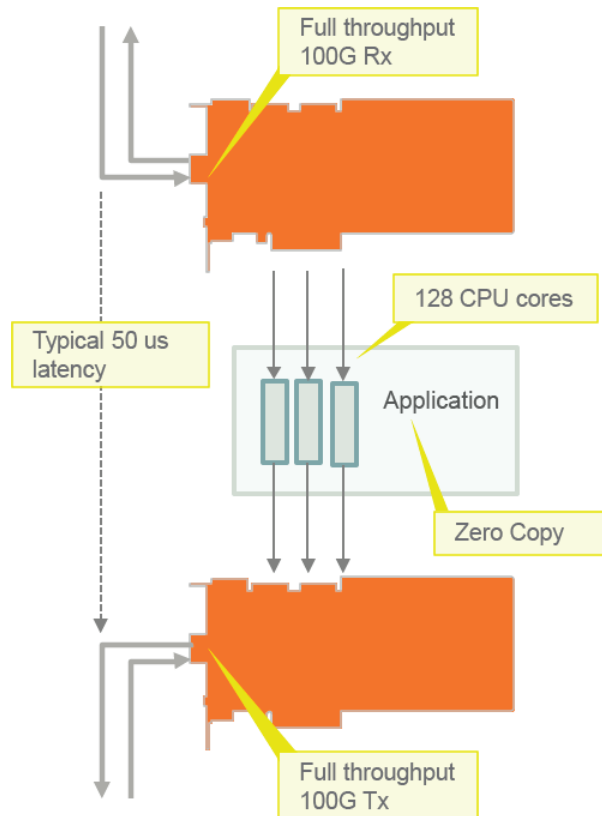
The Napatech 200G performance solution is based on two NT100E3-1-PTP accelerators, each providing a single 100G port with full throughput and zero packet loss under all conditions. Upstream and downstream are combined using QPI Bypass over Intelligent Hardware Interconnect. This allows specific flows to be combined and directed to the right CPU cores for processing in a balanced and controlled manner, without loading the server system with expensive communication between NUMA nodes.

IN-LINE APPLICATION SUPPORT

The Napatech accelerator family supports in-line applications enabling customers to create powerful, yet flexible in-line solutions on standard servers. The more CPU-demanding the application is, and the higher the speeds of links, the higher the value of this solution.

IN-LINE FEATURES

- Full throughput bidirectional Rx/Tx up to 100G link speed for any packet size (Depending on server performance)
- Supports multi-core processing with up to 128 Rx/Tx streams per accelerator
- Customizable hash based load distribution
- Zero copy transfer from Rx to Tx
- Single bit flip to select discard or forward packet
- Typical 50 us roundtrip latency from Rx to Tx fiber



In-Line application support

MONITORING SENSORS

Sensors on the accelerators provide extensive monitoring of:

- 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
- Ethernet link status per port
- Status and loss of time synchronization

The system status can be read from the LEDs on the front of the accelerator. Possible alarms and error codes are stored on the accelerator for later troubleshooting.

HOST-BASED TRANSMIT

Multiple Tx streams enables multi-threaded processing and full throughput bidirectional Rx/Tx. The feature includes an API call for prepending protocol layers to Tx packets.

ADVANCED STATISTICS

RMON1 PORT STATISTICS

Napatech accelerators provide RMON1 (RFC2819) counters on a per-port basis.

ADVANCED PORT STATISTICS

In addition to RMON1 port statistics, the accelerator hardware generates an extensive amount of additional statistics counters, which are available independently of whether the traffic is forwarded to the host or not. This enables customer applications to retrieve a comprehensive network traffic analysis with extremely low CPU load.

Two types of statistics counters are available:

- Large sets: RMON1 (RFC2819) counters with extension of jumbo frame counters are available for both captured and discarded frames on a per-port basis.
- Normal sets: Frame and byte counters for good and bad frames are available per color (filter) and per host buffer.

Counter sets are always delivered as a consistent snapshot time-stamped by a 64-bit high-precision clock and can be synchronized with the Ethernet frames delivered for analysis.

COMPANY PROFILE

Napatech is the world leader in data delivery solutions for network management and security applications. As data volume and complexity grow, organizations must monitor, compile and analyze all the information flowing through their networks. Our products use patented technology to capture and process data at high speed and high volume with guaranteed performance, enabling real-time visibility.

We deliver data faster, more efficiently and on demand for the most advanced enterprise, cloud and government networks. Now and in the future, we enable our customers' applications to be smarter than the networks they need to manage and protect.

Napatech. SMARTER DATA DELIVERY

NAPATECH ACCELERATORS




NT200A01-SCC-2x100



NT100E3-1-PTP



NT80E3-2-PTP



NT200A01-NEBS-2x100



NT100E3-1-PTP-NEBS



NT80E3-2-PTP-NEBS



NT40A01-SCC-4x1



NT20E3-2-PTP




NT200C01-2



NT40A01-NEBS-4x1



NT20E3-2-PTP-NEBS



NT200C01-2-NEBS

**EUROPE, MIDDLE EAST
AND AFRICA**

Napatech A/S
Copenhagen, Denmark

Tel. +45 4596 1500
info@napatech.com
www.napatech.com

NORTH AMERICA

Napatech Inc.
Boston, Massachusetts
Los Altos, California
Washington D.C.

Tel. +1 888 318 8288
info@napatech.com
www.napatech.com

APAC

Napatech China/South Asia
Taipei City, Taiwan
Tel. +886 2 28164533 Ext. 319

Napatech Japan K.K.
Tokyo, Japan
Tel. +81 3 5326 3374

ntapacsales@napatech.com
www.napatech.com