

## Accelerators for Network Analysis

# N-ANL3A 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

#### Time Stamp

- 4 ns time stamping
- Time stamp injection

#### 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
- Deduplication
- Ethernet FCS and IP/UDP/TCP checksum

### Tunneling Support

- GTP and IP-in-IP 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
- Configurable hash keys
- Correlation key

### Frame and Flow Filtering

- Configurable filters
- Port, Protocol and size filters
- IP match filters
- Pattern compare filters
- Coloring

### Intelligent Multi-CPU Distribution

- Configurable data distribution

### Advanced Statistics

- RMON1 port statistics
- Advanced statistics

### Monitoring Sensors

- Temperature and power sensors

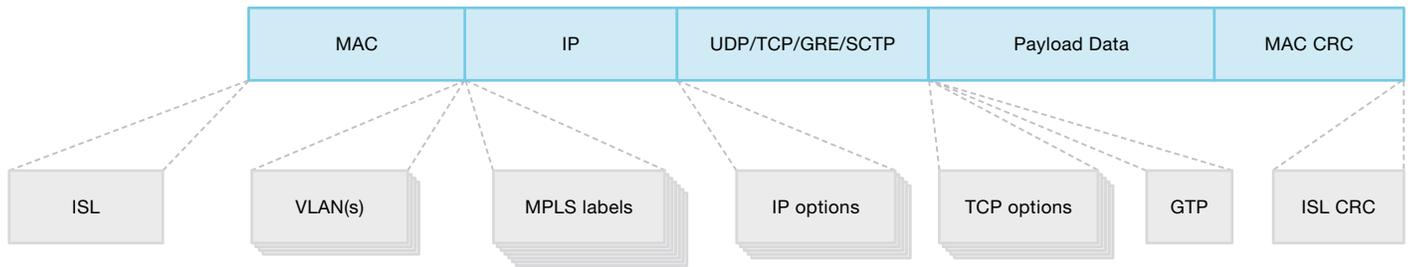
### Managed Transmit

- One-touch frame transmit decision
- Synchronized transmit
- Local retransmit
- Inter-Frame Gap (IFG) control
- Ethernet FCS and IP/UDP/TCP checksum generation

### Programmable Bypass

- Fail-safe using built-in bypass function

## FEATURE DESCRIPTIONS



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.

### 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 4 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 5 different 64-bit time-stamping formats:

- 2 Windows formats with 10 ns or 100 ns resolution
- Native UNIX format with 10 ns resolution
- 2 PCAP formats with 1 ns or 1000 ns resolution

### TIME STAMP INJECTION

A time stamp can be injected into the payload of a packet when transmitting. By supplying an offset, the time stamp can be injected at a convenient point within the packet. As the injected time stamp shows the time the packet was transmitted, it can be used to measure the delay or latency of a network. The receiver of the packet can compare the injected time stamp with the packet time stamp. This shows how long time the packet has been in transit from the transmitter to the receiver.

### 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 either internally, externally or 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.

Onboard IEEE 1588-2008 support enables Napatech accelerator synchronization over longer distances than the 20 m restriction of coax cables – up to 100 meters for direct connections and further over switched networks.

Napatech accelerators with onboard IEEE 1588-2008 support include Default, Telecom (ITU-1 G.8265.1) and Power (IEEE-C37.238) profile support. Accelerators can be used as master or slave in default profile or as slave in Telecom and Power 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, PPPoE, 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, IP, TCP, UDP

## 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
- 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

## DEDUPLICATION

When network traffic is analyzed, large amounts of duplicate frames are sometimes received. Napatech accelerators have hardware functionality that can recognize and remove duplicate frames. This saves a substantial amount of server CPU cycles.

## ETHERNET FCS AND IP/UDP/TCP CHECKSUM

For each received frame, the Ethernet Frame Check Sequence (FCS) is checked, as well as the checksum information in IP, UDP and TCP headers. For in-line applications, the FCS and checksum can be re-used or re-generated as required.

## TUNNELING SUPPORT

### GTP AND IP-IN-IP 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.

### 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.

### 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.

### CONFIGURABLE HASH KEYS

In addition to the built-in protocol hash-keys provided in Napatech Accelerators, it is also possible to configure hash-keys based on selection of fields from encapsulated protocol headers. This allows specific flows to be identified and distribution of CPU processing load to be optimized based on the prevailing traffic.

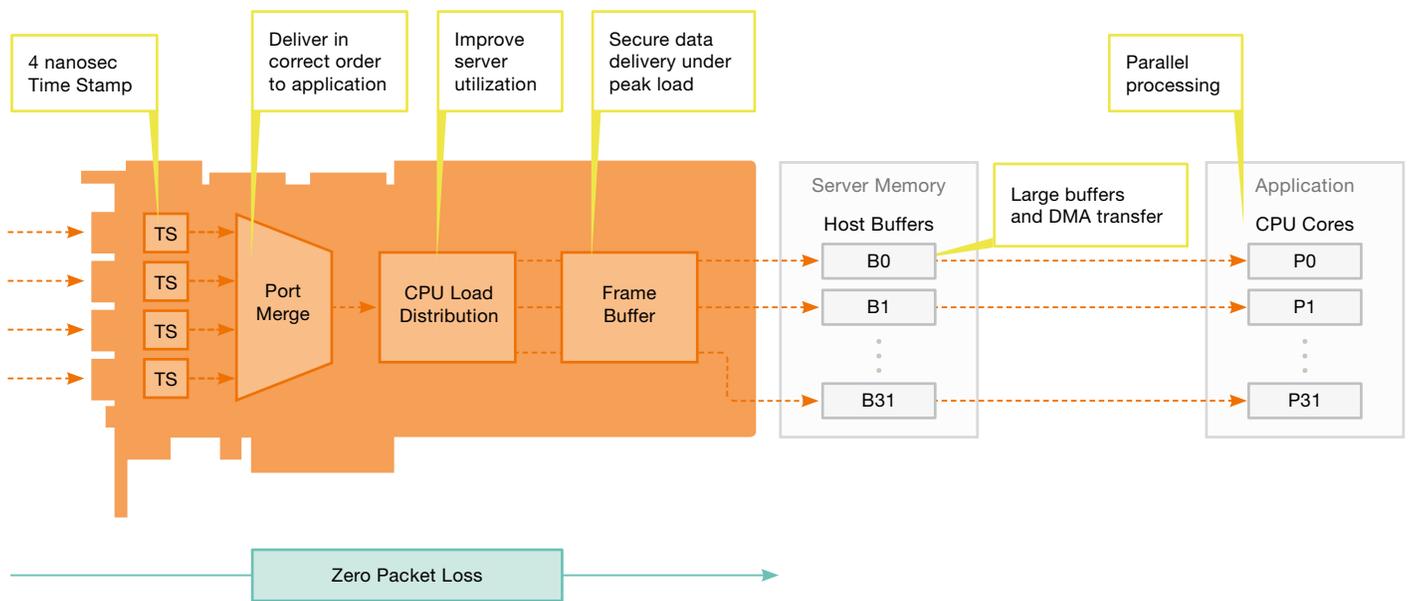
### CORRELATION KEY

With Napatech accelerators it is possible to generate a correlation key that can be used to monitor individual packets at multiple points in the network. The correlation key is a unique identifier for individual packets and can be used as an alternative to IP source and destination addresses for cases where network address translation can change IP addresses in the network being monitored. With correlation keys it is possible to measure the latency at multiple points in the network on a packet-by-packet basis.

### 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.

In all, a single filter will give the correct output under 7,744 different traffic conditions. Napatech accelerators have 64 programmable filter blocks, which can be combined in various ways.

### **PORT, PROTOCOL AND SIZE FILTERS**

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

### **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 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.

### **INTELLIGENT MULTI-CPU DISTRIBUTION**

Multi-CPU buffer splitting enables accelerators to place captured frames in 1 to 32 host buffers. The customer can configure the size of the host buffers from 16 MB to 128 GB; 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.

### **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.

### **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.

### **MANAGED TRANSMIT**

#### **ONE-TOUCH FRAME TRANSMIT DECISION**

Received frames can be quickly and easily transmitted using a single command. A transmission indicator can be set to either transmit the frame immediately or using the received timing.

## **SYNCHRONIZED TRANSMIT**

Transmission or replay is controlled by transmission clocks. Synchronized and coordinated transmission is therefore possible across multiple ports and accelerators when their time stamp clocks are synchronized. Synchronized transmit is also possible over multiple appliances and geographical locations when using time synchronization. This is ideal for traffic generation and capture replay applications.

## **LOCAL RETRANSMIT**

With local retransmit it is possible to automatically re-direct received data flows for re-transmission to one or several ports on the same accelerator. Selection of frames and ports for retransmission is made using filtering logic. It is possible to both retransmit and analyze frames at the same time for applications that need to provide the same data to multiple analysis appliances. In this case, the data for analysis can be selected using filtering logic and forwarded to one or more external appliances.

## **INTER-FRAME GAP (IFG) CONTROL**

For transmitted frames, the timing of transmission can be controlled by adjusting the IFG. Frames can be transmitted with the original IFG or the IFG can be adjusted higher or lower as required.

## **ETHERNET FCS AND IP/UDP/TCP CHECKSUM GENERATION**

For each transmitted frame, the Ethernet Frame Check Sequence (FCS) and checksum information in IP, UDP and TCP headers can be generated automatically by the accelerator.

## **PROGRAMMABLE BYPASS**

For in-line applications, a fail-safe mechanism is required to ensure data traffic continuity in the event of an appliance failure. Napatech provides programmable bypass functionality that allows the user to program different modes of bypass operation, both for power on and power off, including a watchdog timer for automatic failover.

The programmable bypass functionality is provided independently for each port or port pair supported by the accelerator. Three modes are available:

- Normal mode where the port is connected to the accelerator PHY
- Bypass mode where the port pair is interconnected
- Disconnect mode where the port is disconnected

The watchdog timer with software-programmable time-out interval operates independently on each port pair.

## **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. FASTER THAN THE FUTURE**

## NAPATECH ACCELERATORS



NT100E3-1-PTP



NT80E3-2-PTP



NT40E3-4-PTP



NT100E3-1-PTP-NEBS



NT80E3-2-PTP-NEBS



NT40E3-4-PTP-NEBS



NT20E3-2-PTP



NT4E2-4T-BP



NT4E-4-STD



NT20E3-2-PTP-NEBS



NT4E2-4-PTP



NT4E-4-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  
Mountain View, 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

Napatech Korea  
Seoul, South Korea  
Tel. +82 2 6001 3545

ntapacsales@napatech.com  
www.napatech.com