

Mellanox NPA-2™ and NPA-3™



Access Network Processors with Integrated Traffic Management

Mellanox's NPA is a highly-flexible network processor family with integrated traffic management targeting Ethernet network access platforms (ONT/OLT GPON/EPON), copper access platforms (DSLAM) and demarcation devices and 3G/4G WiMAX base stations aggregation and backhaul. The NPA provides high integration, programmable packet processing and advanced flow-based bandwidth control at 10-Gigabit aggregate throughput.

HIGHLIGHTS

- Single-chip, programmable, wire-speed network processor with 10-Gigabit aggregate throughput
- Scaled-down version of Mellanox's NP-3™ network processor targeting Ethernet access applications
 - Software compatible with Mellanox's NP-2™ and NP-3™
- Flexible processing with programmable packet parsing, classifying, modifying and forwarding enabled through integrated Task Optimized Processors (TOPs)
- Integrated 10-Gigabit traffic management with hierarchical scheduling, supporting services defined by the Metro Ethernet Forum, e.g. MEF9 and MEF15
- On-chip Fabric Interface Controller (FIC) functionality for direct interfacing to Ethernet fabrics enabling system-wide traffic management
- Integrated hardware implemented search engines
- Integrated 1MB memory for lookup tables and statistics counters
 - Extension of lookup tables to external DRAM
 - Extension of statistics counters to external DRAM
- Integrated TCAM for on-chip ACL processing
- On-chip OAM protocol processing offload
- On-chip IEEE1588v2 clock sync processing offload
- 1-lane PCI-Express 2.5Gbps external host CPU interface for configuration and control
- RGMII data-only interface to host CPU
- Supports oversubscription beyond 10-Gigabit processing throughput by smart classification

PACKAGE / PROCESS / POWER

- Package: BGA 676 pins, 1.00 mm pitch, 27x27 mm
- Process: TSMC 90nm
- Est. power dissipation: 6-10W typical
- RoHS compliant
- Industrial operating temperature range: -40°C to 85°C ambient

TARGET APPLICATIONS

- Line card, service card and pizza box applications
 - Optical access (GPON/EPON OLTs and ONUs)
 - Wireless backhaul and base station aggregation (3G/4G and WiMax)
 - Access routers
 - Fiber and copper Ethernet access switches
 - Ethernet demarcation devices
 - Copper access (DSLAMs)
- Programming delivers a variety of applications such as L2 switching, Q-in-Q, PBT, T-MPLS, VPLS, MPLS and IPv4/IPv6 routing

MODELS

- Footprint and software compatible models:
 - **NPA-2:** Sixteen 1-Gigabit Ethernet ports (SGMII/SERDES)
 - **NPA-3:** Configurable to eight 1-Gigabit Ethernet ports (SGMII/SERDES) and two 10-Gigabit Ethernet ports (XAUI); or twelve 1GE and one 10GE ports; or four 10GE ports
- SGMII interfaces support 2.5Gbps SERDES for GPON applications

DETAILED FEATURE LIST

Integrated Traffic Management

- 10-Gigabit traffic manager providing queuing and scheduling on all transmitted traffic on all ports
- Per Flow Queuing (PFQ) with 4 level hierarchical scheduling:
 - 32 ports/channels
 - 256 sub-ports
 - 4K classes/subscribers
 - 16K flow queues
- Policing: Per-flow metering, marking and policing
- Configurable WRED profiles
- Shaping: Single and Dual leaky bucket controlling committed/peak rate/bursts (CIR, CIB, PIR, PIB) with IFG (Inter Frame Gap) emulation for accurate rate control
- Scheduling: WFQ and priority scheduling at each hierarchical level
- Work conserving and non-work conserving schedulers
- Frame size from 1 byte to 16K bytes
- Up to 0.5 Gbyte total frame buffering in external DRAM
- Per-frame timestamp and timeout drop
- Dynamic hitless reconfiguration and resource allocation
- LAG shaping
- Out-of-band flow control per physical port (SGMII or XAUI) or logical channel (total of 32 channels)

Programming

- Single-image programming model with no parallel programming or multi-threading
- Automatic allocation of frames to processing engines (TOPs) with passing of messages among TOPs
- Automatic ordering of frames
- In-service software updates
- Large code space memory for multiple and complex applications

Integrated Search Engines

- Performs flexibly defined lookups in switching, routing, classification and policy tables
- Programmable size and contents of search keys and results (associated information) per table
- Support for long keys and long results per table entry
- Table entries stored in integrated memory for fastest lookup time
- Tables may be stored in external DRAM memory
- On-chip state learning and updates of millions of addresses, sessions and flows per second

Integrated TCAM

- Enables fast lookups through tables with wildcards, such as Access Control Lists (ACL)

Statistics and Counters

- Per-flow statistics for programmable events, traffic metering, policing and shaping
- Programmable threshold settings and threshold exceeded notification
- Dynamic allocation and auto association between counters and flows. Counters are automatically recycled when a flow is deleted or aged.
- Hardware implementation of token bucket per flow (srTCM, trTCM or MEF5)

OAM Offload

- Per OAM session state tracking and reporting
- 802.1ag and 802.3ah compliant OAM offload
- Dedicated timer hardware blocks
 - KeepAlive frame generation for precise and accurate session maintenance operations
 - KeepAlive watchdog timers for fastest detection time
- Flexible statistics collection on a per session basis

Integrated FIC Functionality

- For architectures that adapt standard Ethernet switches as the SF solution, the NPA integrates the FIC functionality
- Allows use of standard low-cost Ethernet switches as the backplane switch fabric
- Direct connection from NPA on the line card to the backplane Ethernet switch
- NPA provides for system-wide QoS with per COS and per-flow congestion management

Sync Ethernet

- Enables on-board clock generation schemes using an external or recovered clock reference
- Provides output clock selection from each Serdes lane recovered clock

IEEE1588v2

- On-chip IEEE1588v2 clock sync processing offload for precise time synchronization among remote nodes and switches
- Can operate as the clock master, boundary clock, transparent node or a combination thereof
- Provides an accurate RTC, adjustable from the control CPU or an external source, and provides input and output timestamping for time and delay measurement

SYSTEM CONFIGURATIONS

NPA's flexibility and integration allows system vendors to deliver cost effective solutions that can easily adapt to changing market requirements. Illustrated below are several sample solutions. external source, and provides input and output timestamping for time and delay measurement

Figure 1.
10GE access switch or 10GE inter-carrier demarcation box (NNI) with redundancy or GE demarcation box (UNI)






Figure 2.
10GE inter-carrier demarcation box (NNI)



In the following application, NPA provides switching and QoS to Radio and DSL line cards, and uplinks to aggregation networks:

Figure 3.
Wireless backhaul aggregation switch or DSLAM ADSL/VDSL



In the following application, NPA provides switching and QoS to GPON/EPON links, and uplinks to aggregation networks :

Figure 4.
GPON OLT




Table 1 - Part Numbers and Descriptions

OPN	Description
20775702, 20775701	NPA-2 Access Network Processor with 16 x GE ports. (ROHS)
20775802, 20775801	NPA-3 Access Network Processor with 16/8/12/0 x GE ports and 0/2/1/4 x 10GE ports respectively. (ROHS)