

Access Network Processor with Integrated Traffic Management and Host CPU Core Product Brief

Highlights

- Single-chip, programmable, wire-speed network processor with 5-Gigabit aggregate throughput
- Flexible processing with programmable packet parsing, classifying, modifying and forwarding enabled through integrated Task Optimized Processors (TOPs)
- Integrated 5-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 memory for lookup structures, statistics counters and TM control structures. Optional extension to external DRAM.
- Integrated TCAM for on-chip ACL processing
- On-chip OAM protocol processing offload
- On-chip IEEE1588v2 clock sync processing
- Synchronous Ethernet support, ITU-T G.8261
- On-chip MIPS34Kc core 650MHz for system control
- Supports oversubscription beyond 5-Gigabit processing throughput by smart classification and handling of up to 10Gbps
- Scaled-down version of EZchip's NPA-1 network processor targeting Ethernet access applications
 - Software compatible with EZchip's NPA-1/2/3, NP-2, NP-3 and NP-4 network processors

Package / Process / Power

- Package: FCBGA 484 pins, 23x23 mm, 1.00 mm pitch
- Process: TSMC 65nm
- Power dissipation: 5W typical
- RoHS compliant
- Industrial operating temperature range -40C° to 85C° ambient



Target Applications

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

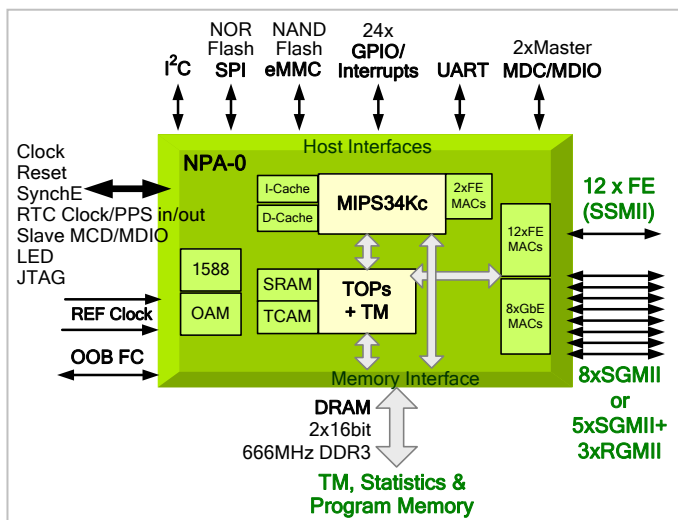
Network I / Os

- 8 x 1-Gigabit Ethernet ports (SGMII/SERDES)
 - 3 ports can be RGMII
- SGMII interfaces support 2.5GE, 3.125Gbps SERDES for GPON
- 12 x Fast-Ethernet ports with SMII interfaces
 - Alternatively 4xSMII + 2xMII, or 8xSMII + 1xMII

Detailed Feature List

Integrated Traffic Management

- 5-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
 - 1K classes/subscribers
 - 256 sub-ports
 - 4K flow queues
- Policing: Per-flow metering, marking and policing
- Hierarchical WRED
- 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 1M frame buffers in external DRAM
- Per-frame timestamp and timeout drop
- Dynamic hitless reconfiguration and resource allocation
- OOB flow control and status interface
 - Elaborated congestion status reports
- Out-of-band flow control per physical port (1GbE or Fast Ethernet) or logical channel (TM level entity)
- Internal buffering and TM external buffering congestion status reporting via OOB signaling

NPU Programming

- Single-image programming model with no parallel programming or multi-threading
- Automatic allocation of frames to processing engines (TOPs)
- SW messages between TOPs stages
- Automatic ordering of frames
- In-service software updates
- Large code space memory for multiple and complex applications
- Microcode compatible with EZchip's NPA-1/2/3, NP-2, NP-3 and NP-4 network processors

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

External Memory

- One interface with two controllers sharing the address and control buses, 2 x 16bit 666 MHz DDR3 DRAM
- Used for TM data & control, external lookup structures, external statistics and MIPS CPU memory
- Flexible bandwidth allocation for various tasks

Statistics and Counters

- Stored in integrated memory (shared with lookup structures) and/or in external DRAM
- 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)

Embedded CPU Core

- 650 MHz MIPS34Kc RISC with 9 stage execution pipeline. MIPS 16e Code Compression
- Power-down mode (WAIT)
- Bootable from ROM/NOR-flash via Serial Peripheral Interface or eMMC
- I²C (Master/Slave) and MDC/MDIO (Master) serial management interfaces
- UART for management console

OAM Hardware Offload

- Per OAM session state tracking and reporting
- 802.1ag, 802.3ah and ITU-T Y.1731 compliant OAM offload
- Dedicated timer hardware blocks
 - KeepAlive (CCM) frame generation (as fast as 3.3ms) for precise and accurate session maintenance operations
 - KeepAlive watchdog timers for fastest detection
- LBM/LBR Message generation and detection
- LTM/LTR Message generation and detection
- Flexible statistics collection on a per session basis

Integrated FIC Functionality

- For architectures that adapt standard Ethernet switches as the switching fabric solution, the NPA-0 integrates the FIC functionality
- Allows use of standard low-cost Ethernet switches as the backplane switch fabric
- Direct connection from NPA-0 on the line card to the backplane Ethernet switch
- 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 ordinary clock, boundary clock, transparent clock, 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
- HW assisted two-step IEEE1588v2 protocol

Ordering Information

- NPA-0 RoHS: P/N 20782900

NPA-0 Architecture

EZchip's NPA-0 is a highly-flexible network processor 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 for aggregation and backhaul. The NPA-0 provides high integration, programmable packet processing and advanced flow-based bandwidth control at 5-Gigabit aggregate throughput.

Through programming, the NPA-0 delivers a variety of applications such as L2 switching, Q-in-Q, MAC-in-MAC, PBB-TE, PBT, T-MPLS, MPLS-TP, VPLS, MPLS and IPv4/IPv6 routing. The integrated traffic management provides advanced QoS for flow-based service level agreements (SLA) and for enabling triple-play services (voice, video, data).

Integrates functions normally found in several chips:

- Programmable packet processing at 5-Gigabit throughput
- Traffic manager
- Classification search engines
- On-chip control CPU 650 MIPS
- OAM processing offload
- Synchronous Ethernet and IEEE 1588v2 clock sync including boundary clock support
- On-chip memory for lookup structures and statistics
- On-chip TCAM for ACLs
- On-chip Ethernet Fabric Interface Controller (FIC)
- Eight Gigabit and twelve Fast Ethernet MACs

NPA-0 provides exceptionally **flexible packet processing** enabling system designers to future proof their designs to support new protocols and features through s/w updates. Packet parsing is supported for any field anywhere in the packet. Various table lookup options are provided with support for long lookup keys and results. Flows are classified based on any combination of extracted packet information. Any packet header and content can be edited and packets can easily be replicated to support multicast applications. A 'run to completion' processing model guarantees support for processing scenarios of any complexity. Large code space is provided to support complex applications as well as true hitless code updates.

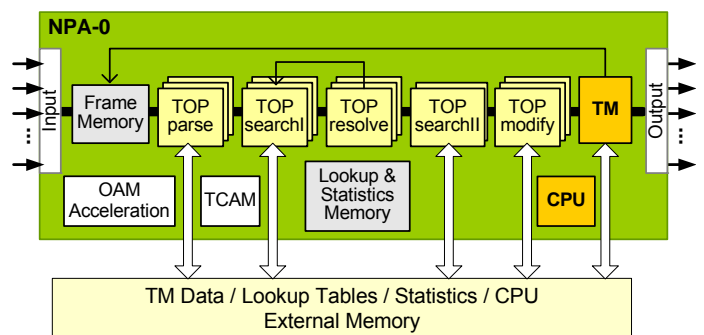


Figure 1. NPA-0 Block Diagram

EZchip's innovative **TOPcore® technology** enables the NPA-0 to deliver its exceptionally high performance. TOPcore technology integrates many high-speed processors; each optimized to perform a specific task. Four types of TOPs (Task Optimized Processors) – parse, search, resolve and modify – are employed to perform the main tasks of packet processing, i.e. classification, forwarding and modification. A programmable TOP corresponds to each of these tasks, and performs its respective task exceptionally fast.

Each TOP processor type employs a unique architecture with a customized, function-specific data path and instruction set. This minimizes the number of clock cycles required for complex packet manipulation and provides exceptionally fast packet processing. TOP performance is boosted by a super-scalar architecture in which multiple instances of the TOPs operate in parallel within each pipeline stage.

NPA-0 uses a simple **single-image programming model** with no parallel programming or multi-threading. Allocation of the TOPs processing engines to incoming frames, passing messages between the TOPs as well as maintaining the ordering of frames is completely transparent to the programmer and performed in hardware. Large code space memory is available to support multiple and complex applications while providing headroom for adding new features. Full support for hitless code updates is provided to reduce system downtime for maintenance.

NPA-0 offers extensive **traffic management** capabilities for traffic transmitted on all NPA-0 interfaces. Flows transmitted to the network links and the control CPU interface can be assigned with specific QoS settings, queued and aggregated to enforce SLAs for services, subscribers, virtual ports and ports.

NPA-0 supports DiffServ and IntServ services and a wide variety of QoS mechanisms. These include:

- Classification – assigning frames to specific flows with applicable QoS parameters.
- Metering – measuring per-flow traffic and determining compliance with traffic parameters. Single and Two Rate Three Color Marking (srTCM and trTCM) are used, compliant with the Metro Ethernet Forum specifications.
- Marking – individual frame's compliance.
- Congestion Avoidance – profile-based WRED early packet discards based on priority, available memory and metering results.
- Traffic Conditioning – enforcing rules on traffic flows, e.g. policing in which packet dropping is applied to non-compliant packets, or shaping which schedules a flow to conform to its assigned parameters (no packet dropping). Single and dual token bucket shaping is applied.
- Congestion Management – hierarchical scheduling of flows to the various interfaces using a priority scheme and Weighted Fair Queuing (WFQ).
- Fabric Interface Chip functionality – Inter device messaging to avoid switch fabric congestion and target output queue congestion.

When using the NPA-0 as a line card device in a chassis, the NPA-0 enables **use of standard low-cost Ethernet switches as the chassis backplane switches**. In this configuration NPA-0 connects directly to the backplane through 1-Gigabit or 2.5-Gigabit serial interfaces, and provides the required QoS features on top of the Ethernet switch features. The NPA's integrated TM queues and buffers manage the traffic transmitted to the backplane Ethernet switch and provide system-wide congestion management per COS and per flow.

NPA-0 contains **integrated memory** for lookup structures and statistics counters. In addition, standard DDR3 SDRAM offers additional external memory. The external DRAM is used also for Traffic Manager frame buffering and control descriptors.

NPA-0 features **integrated search engines** that perform lookups for implementing diverse applications in layer 2-4 switching/routing and layer 5-7 deep packet processing. These search engines deliver programmable lookups in a combination of tables. NPA-0 can store the **lookup structures in integrated memory** to reduce board complexity, power dissipation and cost. Memory is available on-chip and applications can also utilize the external DRAM. The integrated memory enables numerous lookups to be performed per packet while sustaining high throughput.

NPA-0 supports three types of **lookup structures** – direct access tables, hash tables and FastIP – each is flexibly defined and used for various applications. Tables may be used for forwarding and routing, flow classification, access control, etc. Numerous tables of each type can be defined, stored in internal memory and/or external memory and searched through per packet. For maximum flexibility, the key size, result (i.e. associated data) size, and number of entries are all user-programmed per table.

Search algorithms for hash and FastIP are implemented in hardware for maximum performance and simplicity. Hash lookup performance is nearly deterministic regardless of the hash table size. FastIP is an innovative implementation of the routing table as a series of direct access tables as opposed to a tree. The FastIP data structure improves lookup performance and reduces the number of memory accesses required, making it ideally suited for IPv4 routing tables. FastIP can be used to perform best match IP address longest-prefix-match lookups. Up to 96 bytes of result (associated information) can be stored per each entry and retrieved upon a lookup match.

In addition, an **integrated TCAM** is available for performing fast lookups through tables with multiple wildcards such as Access Control Lists (ACL). The TCAM lookups are performed in parallel to algorithmic lookups performed by the integrated TOPsearch engines of NPA-0.

Stateful auto learning and updating of table entries and the session's state is performed entirely by NPA's TOPs with no intervention required by the CPU. TOPs hardware delivers a high learning rate of millions of addresses per second, for the addition of new flows or deletion of old flows to/from the forwarding database or flow table. Result information, statistics counters, and per-flow rate limiters can be automatically associated with new flows and recycled when deleting old flows.

On-chip **OAM support** tracks individual sessions and offloads the host CPU from the task of generating and monitoring OAM messages. Dedicated configurable timers generate KeepAlive frames for thousands of sessions. These frames are processed by the TOPs for flexible formatting and accounting on a per session basis. Dedicated h/w monitors thousands of sessions and verifies that KeepAlive messages arrive within configurable intervals and according to a specified rate. The host CPU is alerted to any session failing to meet the minimum messaging rate.

NPA-0 provides support for **Synchronous Ethernet** and the **IEEE1588v2** protocol for precise time synchronization among remote nodes and switches. For Synchronous Ethernet the NPA-0 enables on-board clock generation schemes using an external or recovered clock reference. It provides clock recovery per Serdes lane and configurable output clocks for each Serdes lane. NPA-0 can operate as the boundary clock, transparent node or combination of these. NPA-0 provides an accurate real time clock, adjustable from the control CPU or an external source, and provides input and output time stamping for time and delay measurements. Time synchronization is assisted with the on chip OAM block for periodical synchronization transactions without requiring CPU intervention, to enable data plane implementation of IEEE1588 protocol on-chip.

For application scenarios that load the NPA-0 with traffic bursts that exceed 5Gbps, the NPA-0 provides **smart oversubscription** to assign preferential processing for high priority traffic. Basic traffic parsing and classification are provided for 10Gbps rates. At these rates packets can be classified according to VLAN, priority bits, selected list of SA or DA, MPLS tag and more. The NPA-0 compares the status of its input queues and if a threshold is exceeded, a drop decision can be taken for packets classified as lower priority.

System Configurations

NPA-0'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.




Figure 2. Access switch or inter-carrier demarcation box (NNI) with redundancy




Figure 3. 1GE inter-carrier demarcation box (NNI)

In the following applications, NPA-0 provides switching and QoS to Radio and DSL line cards, and uplinks to aggregation networks:

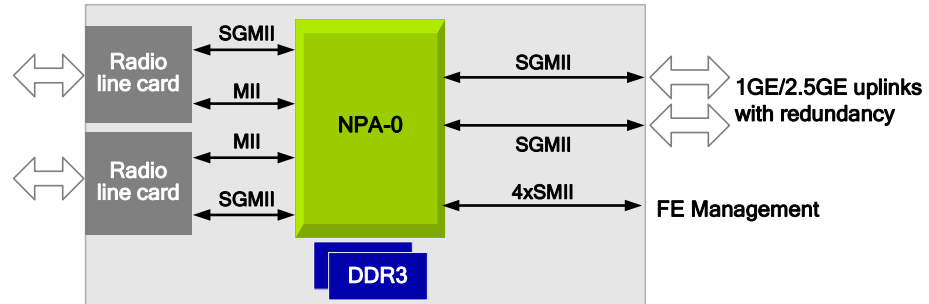


Figure 4. Wireless backhaul aggregation switch




Figure 5. Wireless backhaul CES Gateway

In the following application, NPA-0 provides switching and QoS to EPON/GPON links, and uplinks to aggregation networks:




Figure 6. EPON/GPON OLT

EZdesign Software Toolset and Libraries

EZdesign is a comprehensive set of design and testing software tools for developers, enabling rapid delivery of new designs based on EZchip's network processors. EZdesign allows designers to create, verify and implement NPA-0 applications to meet specific functionality and performance targets.

EZdesign components include:

- **Microcode Development Environment:** A unified GUI for editing and debugging code, including setting breakpoints, single-stepping program execution and access to internal resources. Features include a code editor, view of memory and register contents, performance charting, macro recording and script execution. The EZdesign MDE is used in development and debugging of code on both the simulator and the actual network processor.
- **Simulator:** Provides cycle accurate simulation of the EZchip network processor for code functionality testing and performance optimization.
- **Assembler and Preprocessor:** Generating optimized code for execution on EZchip's network processors. The assembly is interleaved with high-level macros.
- **Applications Library:** Sample code implementing high-level applications for reference when designing new networking platforms and services. Sample code is available for L2 switching, Metro Ethernet, OAM, VPLS, MPLS LER and LSR, IPv4 and IPv6 routing, Network Address Translation (NAT), Access Control Lists (ACL), firewall, server load balancing and more.

- **Frame Generator:** A GUI guiding the programmer through the process of creating frames, layer by layer. Allows for generation of frames of different types, protocols and user-defined fields.
- **EZconfig:** A GUI enabling configuration of the EZchip network processor and definition of data structures used by the network processor for forwarding and policy table lookups (e.g. hash), their keys and associated result information.
- **Traffic Manager Configurator:** A high level GUI enabling you to configure both of the traffic managers embedded in the NPA-0 network processor. This includes configuration of the hierarchical topology and QoS parameters for each level. A configuration sanity check and conversion of NPSl (network processor script language) to C source files are provided.

EZdriver Control Processor API Layer

EZdriver SDK is a toolset that facilitates the development of the control path software for EZchip network processor based systems. It enables applications that run on the control CPU to communicate with the EZchip network processor. EZdriver consists of routines that execute on the control CPU and provide an API for interfacing to the network processor. It includes the chip configuration, microcode loading, creation and maintenance of lookup structures, sending and receiving frames to and from the network processor, as well as configuration of and access to the statistics block.

EZdriver along with EZdesign provide extensive debugging capabilities, and enable software-driven debugging features (e.g. breakpoints, single step, register and memory access) to be performed on both the simulator and the actual network processor.

About EZchip

EZchip Technologies is a fabless semiconductor company that provides Ethernet network processors. EZchip provides its customers with solutions that scale from 1-Gigabit to 200-Gigabits per second with a common architecture and software across all products. EZchip's network processors provide the flexibility and integration that enable triple-play data, voice and video services in systems that make up the new Carrier Ethernet networks. Flexibility and integration make EZchip's solutions ideal for building systems for a wide range of applications in telecom networks, enterprise backbones and data centers. Visit our web site at www.ezchip.com.



Email: ezsupport@ezchip.com • Web: www.ezchip.com

EZchip Technologies Inc. • 900 E Hamilton Ave, Suite 100, Campbell, CA 95008, USA • Tel: (408) 879-7355, Fax: (408) 879-7357
EZchip Technologies Ltd. • 1 Hatamar Street, PO Box 527, Yokneam 20692, Israel • Tel: +972-4-959-6666, Fax: +972-4-959-4166

©2011 EZchip Technologies. All rights reserved. Information is subject to change without notice. EZchip is a trademark of EZchip Technologies. I2C is a registered trademark Philips Electronics N.V. Other brand and product names are trademarks or registered trademarks of their respective holders. Revised; Nov. 20, 2011.