

Mellanox Indigo™ NPS-400™ Network Processor

400Gb/s NPU for Smart Networks

A game changing network processor that merges NPU performance with CPU flexibility and ease of programmability. The Indigo NPS-400 C-programmable packet processor at ultra-high 400-Gigabit throughput, supports all 7-layers with integrated traffic management to enable the next wave of high-performance intelligent carrier routers, data-center network equipment, and accelerates virtualized functions in emerging SDN and NFV networks.

INNOVATIVE CTOPs

Indigo provides C-programming capabilities and a standard Linux OS for both data and control plane. CTOPs (C programmable Task Optimized Processors) offer the benefit of a C based programming environment, while retaining all the benefits of key TOP specialized instructions, which are accessed thru in-line functions. Unlimited code space for control plane and debugging features, a Linux environment that offers access to all the standard toolchain, debugger and other features familiar to programmers.

The NPU architecture is built on an array of symmetrical processors, CTOPs, which run software in a run to completion model. The core hardware multi-threading conceals the latency of memory accesses. The EZ-ISA instruction set, which includes legacy TOP instructions, provides a smooth transition from the NP-5 TOP to the Indigo CTOP architecture.

INDIGO PROGRAMMING

Mellanox's Software Development Kit (EZdk) is a comprehensive set of design and runtime tools for developers, enabling a short time-to-market for new designs based on Mellanox's C programmable Indigo. EZdk allows developers to design, implement and verify both control-plane and data-plane applications for Indigo. EZdk provides an Eclipse™-based IDE (Integrated Development Environment), optimized GNU toolchain and Linux kernel, control and data plane software development kits, reference applications, and SW simulator.

Development Tools

- EZide – Eclipse™-based Integrated Development Environment
- Industry standard-based symbolic software debugging and profiling tools (GDB, perf) with multi-core support
- System on a chip (SoC) level performance profiling and hot-spot analysis tools, allowing detection of system-wide bottlenecks (DDR, etc.)
- Hardware level debugging via JTAG interface
- GUI based editor for generating Indigo device configurations
- GUI based interface for exploring and controlling Indigo device status
- Enables operation with the real hardware targets and the software simulation tools

HIGHLIGHTS

FEATURES

- C programmable 400Gbps network processor for smart networks
- Based on innovative C programmable Task Optimized Processors (CTOPs)
- NPU with 256 CTOPs, 4K threads, traffic manager, DPI, security, search engine and 960Gbps I/O in one chip
- 400-Gigabit wire-speed packet processing
- Advanced services and full layer 2-7 stateful processing
- Easy to use C-programmable and Linux® operating system
- Addressing carrier, cloud and data-center networking equipment

TARGET APPLICATIONS

- Carrier Ethernet switch/router line cards and services cards
- Data center and carrier appliances
- Mobile packet network infrastructure – evolved packet core (EPC) routers
- Software defined networking (SDN), OpenFlow and virtual networking
- Load balancing
- Firewall and intrusion detection and prevention systems (IDS/IPS)
- IPsec and SSL VPN gateways
- TCP offload
- Traffic analytics and network monitoring
- WAN optimization
- DPI and lawful interception
- Session border controllers

- Supported on x86-64 development stations running Linux
- EZsim – Instruction accurate simulation of the CTOPs and the Indigo SoC supporting the running of full multi-core data plane and control plane applications, including SMP (Symmetric Multi-Processing) Linux

Operating System

- SMP Linux 4.x kernel ported to Indigo architecture to run on the CTOPs
- Zero Overhead Linux™ (ZOL) supporting bare-metal performance with the same familiar and convenient Linux user-space development environment
- Linux file system and standard runtime libraries
- Device drivers for the Indigo hardware peripherals, such as the Ethernet network interfaces and the UART interface for console services
- Boot loader (U-Boot)
- Industry standard GNU toolchain

Data Plane Environment Libraries

- Data-plane services libraries providing a C-based application programming interface (API) for data-plane applications running on Indigo, abstracting the complexities of the underlying CTOP core instruction set and various hardware accelerators.

Control Plane Environment Libraries

- Control-plane services libraries providing a C-based API for control-plane applications for Indigo, abstracting the complexities of the underlying hardware interface. The control-plane service libraries are used for configuration and management of Indigo, including its network and memory interfaces, embedded traffic manager and statistics acceleration block, as well as for creation and management of lookup structures
- Access control list compiler for Mellanox's algorithmic TCAM accelerators
- PCRE (Perl Compatible Regular Expression) rule compiler for Mellanox's DPI (Deep Packet Inspection) accelerators
- The control plane environment libraries may be run on an external control CPU or on one or more of Indigo's HW threads implementing the control CPU functionality, for fully autonomous operation without requiring an external control CPU.

Zero Overhead Linux™ (ZOL) Environment

Mellanox provides an integrated SMP Linux runtime environment for both control and data plane applications running on the Indigo CTOP cores.

To achieve the performance requirements for data plane applications, Mellanox's optimized Linux kernel provides deterministic performance for data plane applications on par with a bare metal application, while enjoying the benefits of a standard OS support.

The data plane application processes run on dedicated hardware threads, and are isolated from all uncontrolled interferences. They do not perform context switches unless explicitly requested by the application or the debugger, and are isolated from all interrupt sources, including timer interrupts. The data plane applications access all performance critical code and data through a fixed memory mapping table, avoiding any costly translation lookaside buffer (TLB) misses. The data plane applications interface directly with the hardware services through the data plane user-space library, removing the need for context switching between user and kernel space modes.

The data plane applications are developed, controlled and debugged as standard Linux processes. They can be started/ killed thru a Linux shell, viewed using the standard "ps" or "top" commands and debugged via the standard GDB GNU debugger. Furthermore, the applications can perform any standard Linux library and system call (e.g., "printf"). The Linux kernel still provides the initialization, exception handling and debug services for the data plane applications, but will not involuntarily interrupt these applications unless there is an exception (e.g. divide by zero) or when requested by a debug process.

SOFTWARE LIBRARIES

Mellanox provides a rich set of software middleware and reference applications. Middleware is comprised of reusable building blocks, modules or stacks that build on the basic data and control plane API to facilitate the fast and simple creation of applications.

- Stateful Flow Library
- Deep Packet Inspection Library and DPI-based Application Recognition

Reference Applications:

- Layer 2-3 Quality of Service
- IPsec Gateway
- Layer 4-7 Network Visibility
- Accelerated Linux IP Virtual Server Load Balancer
- Accelerated Data Plane

For more information refer to the *Indigo Software Libraries Product Brief*.

EZdk Software Development Kit

DEVELOPMENT TOOLS

EZide
Eclipse-based
IDE

EZsim
SoC SW
simulator

REFERENCE APPLICATIONS

- Layer 2-3 Quality of Service
- IPsec Gateway
- Layer 4-7 Network Visibility
- Accelerated Linux IP Virtual Server Load Balancer (ALVS)
- Accelerated Data Plane (ADP)

MIDDLEWARE

- Stateful Flow Table (SFT) Library
- Deep Packet Inspection (DPI) Library

DATA PLANE ENVIRONMENT

APIs for data plane services

CONTROL PLANE ENVIRONMENT

APIs for control plane services

OS & RUNTIME ENVIRONMENT

- SMP Linux 4.x
- Linux file system
- Boot loader (U-Boot)
- Programming tools (GNU toolchain)

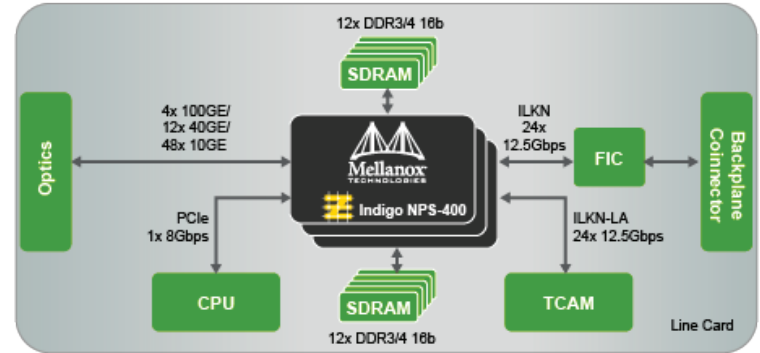
REFERENCE DESIGN

Mellanox also offers compact Indigo hardware platforms which can be used to accelerate the development of target hardware and software applications.

Sample Applications

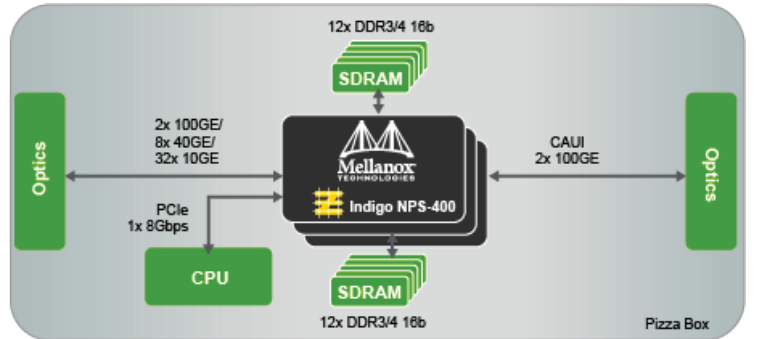
L2-L7 Line Card with FIC, TCAM and Host CPU

This figure shows an over-subscribed line card. The connection to the switch fabric is provided by an Interlaken 24-lane interface via a FIC. An external TCAM for applications requiring large ACL tables is connected using an Interlaken LookAside 24-lane interface. An external host CPU is connected for control applications that require low memory latencies. The CPU is connected with an 8Gbps 1-lane PCIe-host Gen3 port that enables the CPU to access the NPU control registers after boot.



Pizza Box Appliance

In this figure, Indigo provides the entire network layer 2-7 processing and traffic management. Indigo offloads these functions from the CPU. The CPU receives selected packets from Indigo for exceptions handling, and performs all the protocol, management and application tasks. The CPU is connected with a 1-lane PCIe Gen3 port providing 8Gbps throughput. PCIe enables the CPU to access the NPU control registers after boot.



System Application for Multi-Tbps Operation

This figure exemplifies a multi-Terabit/second system with multiple 480Gbps line cards inter-connected through fabric switch cards.

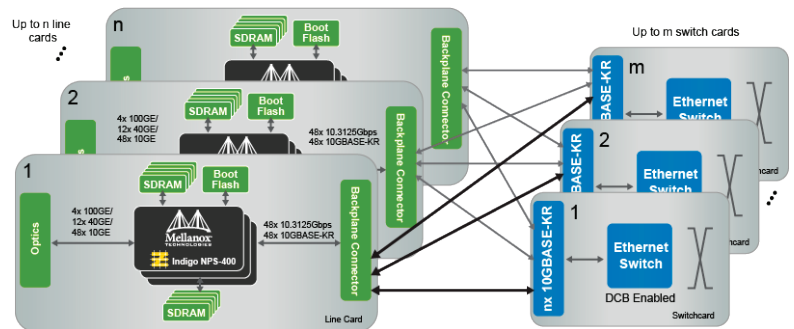
The system provides:

- Multi-Tbps of non-blocking, QoS, unicast and multicast operation
- Multiple Ethernet switches
- Multiple Indigo NPS-400 devices per line card
- Multi-terabit chassis for mega data centers, enterprise and service provider applications

The fabric switches support Data Center Bridging (DCB).

The DCB requires:

- IEEE 802.1Qbb – Priority based flow control (PFC). Indigo supports PFC on all Ethernet links.
- IEEE 802.1Qau – (Quantized) Congestion Notification (CN and QCN). Indigo NPS-400 application supports the standard form and E2E CN proprietary messages.
- 802.1Qaz – Enhanced Transmission Selection by runtime configuration of TM via one of the CTOPs.



Typical number of switch cards (m) is 4 (with 100GbE), 12 (with 40GbE), or 48 (with 10GbE).

Each Indigo device allocates 48x 10GBASE-KR lanes for fabric connectivity. The line card distributes the fabric traffic for even load distribution over the backplane fabric lanes. Each line card maintains virtual output queuing (VoQ) per system output port CoS. In-band signaling of end-to-end QCN messages maintains the CoS aware fabric traffic engineering scheme.

Table 1 - Part Numbers and Descriptions

OPN	Description	OPN	Description
20785001	Indigo NPS-400 400G NPU with traffic mgmt, crypto, RoHS, step 01	20785701	Indigo NPS-400 400G NPU with traffic mgmt, no crypto, RoHS, step 01
20785000	Indigo NPS-400 400G NPU with traffic mgmt, crypto, RoHS, step 00	20785700	Indigo NPS-400 400G NPU with traffic mgmt, no crypto, RoHS, step 00