

NFP-4000 Intelligent Ethernet Controller Family

PURPOSE-BUILT NETWORK FLOW PROCESSOR (NFP) DEVICES FOR INTELLIGENT SERVER ADAPTERS

The cloud data center is experiencing a period of unprecedented growth as it is able to demonstrate significant benefits over a traditional custom infrastructure. Additionally, it is also driven in part by the evolution of data center availability of high-performance servers, virtualization capabilities, need for comprehensive security and ability to deploy new services at a rapid pace.

The challenges associated with supporting new standards, protocols, policies and services have forced the implementations to be driven by software-based approaches. As a result, this has increased the performance requirements imposed on the server CPUs to a point where they are unable to keep up with the processing demands. Therefore, data center operators are using more servers, leading to inefficiencies and increases in capital expenditure and energy consumption.

Netronome's NFP-4000 family of flow processors and associated software provides a practical solution to the challenges encountered when forcing server CPUs to offload expensive tasks and keep alignment with the overall software-driven architecture. The NFP device family is purpose-built to satisfy performance and scaling challenges, and delivers on the promise of intelligent offload-capable intelligent server adapters for data center servers.

Scalable Ethernet Network Interfaces

The NFP-4000 device includes native support for standard x10G, x40G and x100G interfaces. The interfaces are natively capable of auto-negotiation with the north-bound platforms and can directly interface to standard optical, as well as direct attach copper transceivers. The interfaces support datacenter bridging (DCB) requirements, including Priority-based Flow Control (PFC) and Enhanced Transmission Selection (ETS) features.

Multi-Socket Scaling

Netronome delivered the industry's first multi-socket capability; and the NFP-4000 device inherits the legacy by supporting direct simultaneous connectivity to two CPU sockets over PCIe 3.0. Each of the PCIe 3.0 interfaces is capable of supporting up to eight lanes simultaneously. Hardware capabilities for tasks like to load balancing, DMA, queuing and SR-IOV are natively supported by both the NFP-4000 PCIe 3.0 interfaces, resulting in applications being able to perform at line rate.



HIGHLIGHTS

- Intelligent Ethernet controller for server adapter offload applications
- 10x10Gbe, 2x40Gbe, 100GbE Network I/O support
- Single or dual PCIe Gen3x8 connectivity to host CPU
- Software defined data path with standard features supplied by Netronome
- High-performance external memory expansion capabilities for large flow tables

TARGET APPLICATIONS

- Intelligent Server Adapters for cloud data centers
- Open vSwitch acceleration and offload
- Network virtualization with SDN controlled tunnel overlays
- Microsegmentation of security policies
- SDN controlled network monitoring and telemetry

FEATURES & BENEFITS

- Support for open source and standard virtual switch datapath in firmware enables easy integration with hypervisors and higher layer management and orchestration software
- Intelligent data plane processing at higher performance, with lowest power and cost, while enabling capability to adapt quickly to new networking features
- Optimized match action-based flow processing architecture enables per tenant, per VM, per application policies at scale to support micro-segmentation
- Flexible and comprehensive tunneling support enables network virtualization with high performance and scale
- Built-in high-performance bulk cryptography acceleration delivers secure data access at cloud scale

Virtual Switching Acceleration and Offload

Server-based virtual switches such as the ubiquitous Open vSwitch (OVS) are being widely deployed in the data center. The conventional approach of supporting OVS entirely in software on the server CPU suffers from the drawback of requiring a large amount of processing resources that could be otherwise allocated for application processing. The NFP-4000 provides the ability to offload the entire OVS datapath and, at the same, is capable of seamlessly interfacing with the Linux operating system. The NFP-4000 is capable of supporting millions of policy rules using external memory, all at line rate.

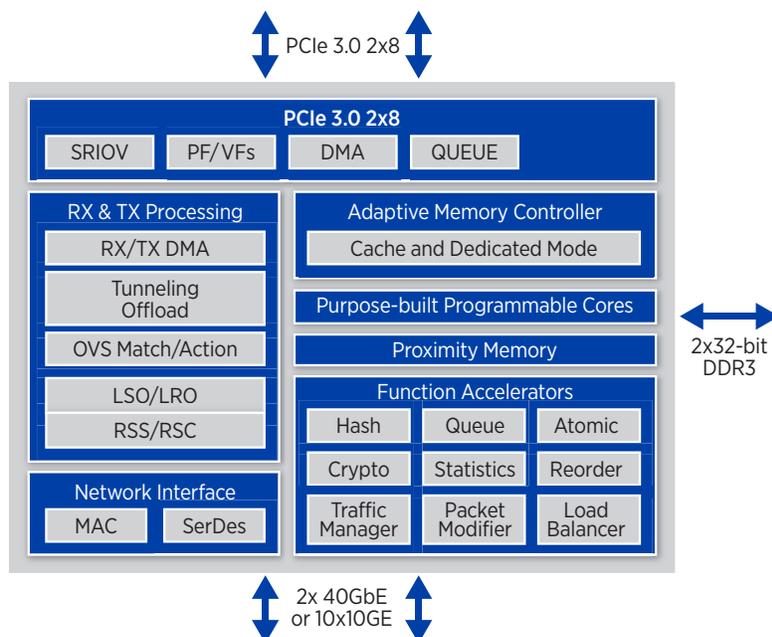
Overlay Tunnel Processing

Tunneling provides the data center operators with a mechanism to virtualize their networks with overlay protocols that terminates into a VM. The tunneled packets transparently pass through the intermediate network elements, like switches. The NFP-4000 with provided firmware natively supports complex encapsulation and de-encapsulation tasks. This allows the datacenter operators to rapidly deploy traditional and emerging protocols such as VXLAN, MPLS, QinQ, and GRE.

Bulk Crypto Security

The data center is a highly virtualized environment that addresses a diverse set of user needs. This includes anytime, anywhere access to personal data and reliance on rapidly deployed new services. Security concerns are consistently identified as a major barrier to this data center transformation. Multiple bulk crypto engines, along with companion processing elements, allow the NFP-4000 to support security protocols such as IPSec that are typically expensive to perform with a traditional host-based solutions.

NFP-4000 Block Diagram



Hierarchical Memory

The large internal memory supported by the NFP-4000 can be flexibly partitioned for storing headers, payloads of packets, exact match, wildcard flow tables, security associations and other data structures, as demanded by various applications. The memory elements themselves are distributed in a hierarchical fashion to deliver highest performance for data center-specific workloads.

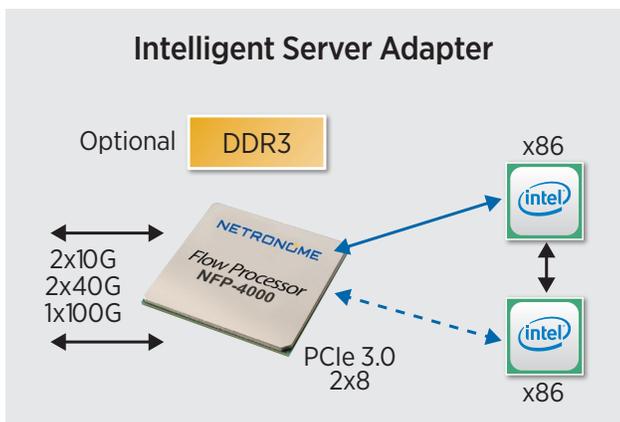
External Memory Controller

The NFP-4000 is capable of supporting up to 12GB of external DDR-3 memory. To overcome penalties associated with DDR-3 accesses, the memory controller incorporates a large amount of internal memory that can adapt as a cache to store the most popular entries entirely via software control. The DDR-3 could also be used for packet storage.

Data Path Programming

With its comprehensive and proven development tools, Netronome supports multiple programming models for server-based networking applications. Flow configuration and Packet I/O API-based programming is supported using Netronome’s Agilio Software. Alternatively, OVSDB or OVS CLI style programming is supported. For programming the networking data path for new features two models are supported: the first is P4-based and is the simplest, allowing for high-level match-action syntax-based hardware-agnostic programming. The second model, allows extensions to the Agilio-SW supported data path using a sandbox extension method that supports both C and P4 based programming. For more details, please see www.Open-NFP.org/resources.

NETRONOME’S NFP-4000 FAMILY OF FLOW PROCESSORS AND ASSOCIATED SOFTWARE PROVIDES A PRACTICAL SOLUTION TO THE CHALLENGES ENCOUNTERED WHEN FORCING SERVER CPUS TO OFFLOAD EXPENSIVE TASKS AND AT THE SAME TIME, KEEPS ALIGNMENT WITH THE OVERALL SOFTWARE-DRIVEN ARCHITECTURE.



	4020-ON-AO-XX	4020-OC-AO-XX	4040-ON-AO-XX	4040-OC-AO-XX	4100-ON-AO-XX	4100-OC-AO-XX
IO (Gbps)	2x10	2x10	4x10, 1x40	4x10, 1x40	10x10, 2x40	10x10, 2x40
PCIe	2x8	2x8	2x8	2x8	2x8	2x8
DDR3	2x40	2x40	2x40	2x40	2x40	2x40
CRYPTO	No	Yes	No	Yes	No	Yes



SPECIFICATIONS

Ethernet	<ul style="list-style-type: none"> IEEE Std 802.3ae 10 Gigabit Ethernet IEEE Std 802.3ba 40 Gigabit Ethernet IEEE Std 802.3ad link aggregation and failover IEEE Std 802.1Q, .1p VLAN tags and priority IEEE P802.1Qaz D0.2 ETS IEEE P802.1Qbb D1.0 priority-based flow control IEEE 1588 precision clock synchronization Jumbo frame support (9.6KB)
PCI Express Interface	<ul style="list-style-type: none"> PCIe 3.0 2x8 Simultaneous dual-socket connectivity Auto-negotiates to x8, x4, x2 or x1
Connectivity	<ul style="list-style-type: none"> Integrated MAC and SerDes for 10/40/100GbE connectivity Native support for SFP+, QSFP, and CXP connectors
Network Acceleration and Offloads	<ul style="list-style-type: none"> TCP/UDP/IP stateless offload Intelligent interrupt coalescence Receive-side scaling Single-Root IOV up to 64 virtual functions per PCIe Link aggregation with LACP DPDK, zero copy, kernel bypass, packet direct
Virtual Switch Data Plane Offload	<ul style="list-style-type: none"> Open vSwitch acceleration and offload Tunnels encapsulation and de-encapsulation, more than 100K tenants Support for custom tunnel types: VXLAN, NVGRE, MPLS Millions of exact match and wild card flows Flow tracking in hardware for cut-through acceleration Stateless and stateful load balancing
Package and Other Information	<ul style="list-style-type: none"> 33mm x 33mm flip chip ball grid array 1mm solder ball pitch 6/6 RoHS-compliant solution

NETRONOME

Netronome Systems, Inc.
 2903 Bunker Hill Lane, Suite 150
 Santa Clara, CA 95054
 Tel: 408.496.0022 | Fax: 408.586.0002
www.netronome.com

©2016 Netronome. All rights reserved.

Netronome, the Netronome logo, and Agilio are trademarks or registered trademarks of Netronome Systems, Inc. All other trademarks mentioned are registered trademarks or trademarks of their respective owners in the United States and other countries.