



Heterogeneous Multicore

Making multicore CPUs work in high-performance designs

Kurt Parker



Agenda

- High-Performance Dataplane Applications
- Engineering Challenges
- Heterogeneous Multicore Architecture
 - Effective Computation
 - Efficient Memory
 - Processor Core Efficiency
- Case Study: Snort
- Avoiding Design Pitfalls
- Network Flow Processors
- Summary and Conclusions

"Many security appliances are adopting multicore x86 processors such as Intel's Nehalem architecture. As the performance requirements for these appliances grow beyond a few gigabits, Netronome's NFP32xx offers performance scaling through data-plane acceleration."

*Bob Wheeler, Linley Group
Security Report, July 2009*

NFP + Multicore x86 > Multicore

Intelligent Dataplane Applications

At 10 Gbps and Beyond

Scaling application performance

Examples: IDS, IPS, UTM, Firewall, Load Balancers, DLP, Forensics, Test & Measurement, Wireless Control, IPTV Ad Insertion, others.

Appliances

Chassis



Horizontal Scaling

- Single software application - low to high end
- Maximize real-world performance vs. marketing

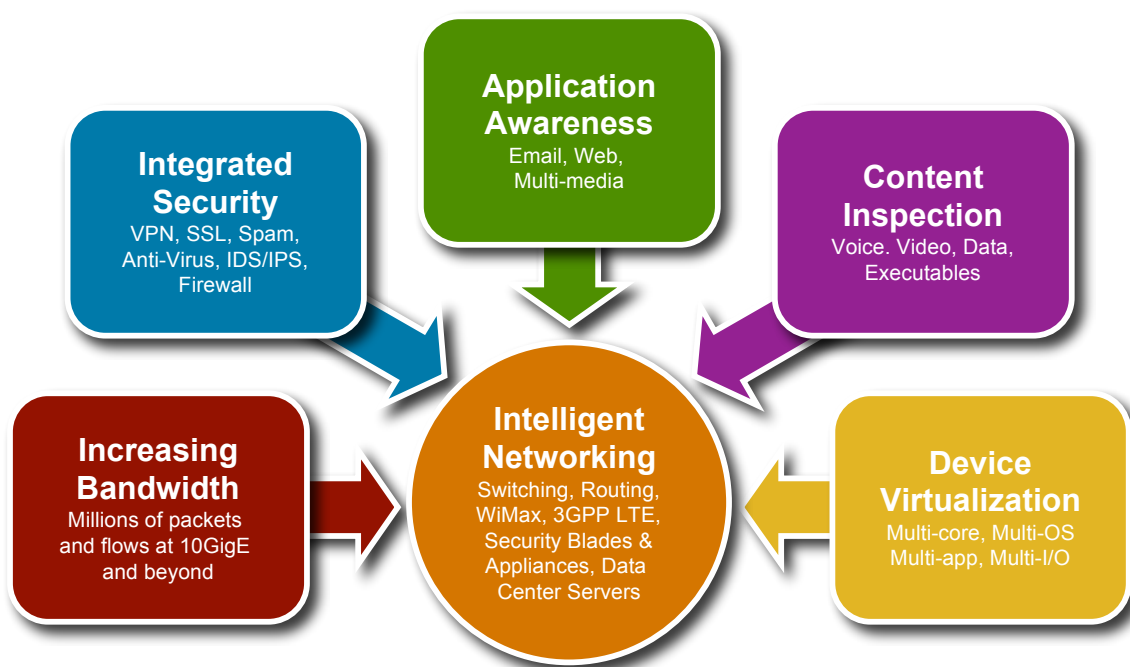
Vertical Scaling

- Increase performance in all systems 2x-4x

Common Characteristics

- Application involved in packet forwarding
- Often far more than L2-L3 forwarding
- Typically computationally intense
- Majority are written to IA/x86 platforms

Design Requirements and Challenges



Source: Morgan Stanley

1 Effective Computation

- Instructions per packet
- DPI, Security, RegEx > 10 Gbps
- Host workload offload

2 Efficient Memory

- Memory bandwidth vs. data rates
- Cache effectiveness


3 Core Efficiency

- Mbps per Watt
- Mbps per mm²

Design Constraints

- High data rates
- Flow-based vs. packet-based
- Non-uniform, multiplexed data impacts cache
- Low latency for interactive inline applications

Heterogeneous Multicore Architecture

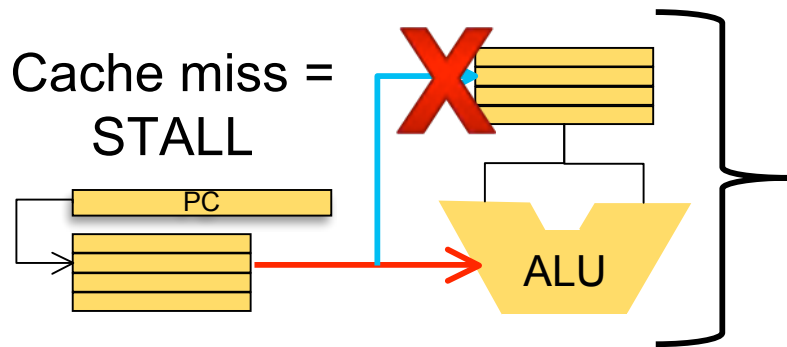


Heterogeneous multicore processing solutions are more than 4x as efficient than x86 and standard NICs

Inter-chip access

- Demultiplexing and classification
- TCP offload
- Host offload for burdensome I/O, security, DPI functions
- Zero copy, big block transfers to multiple cores, VMs or endpoints
- Full I/O virtualization with Intel VTd
- Programmable egress traffic management

1: Effective Computation



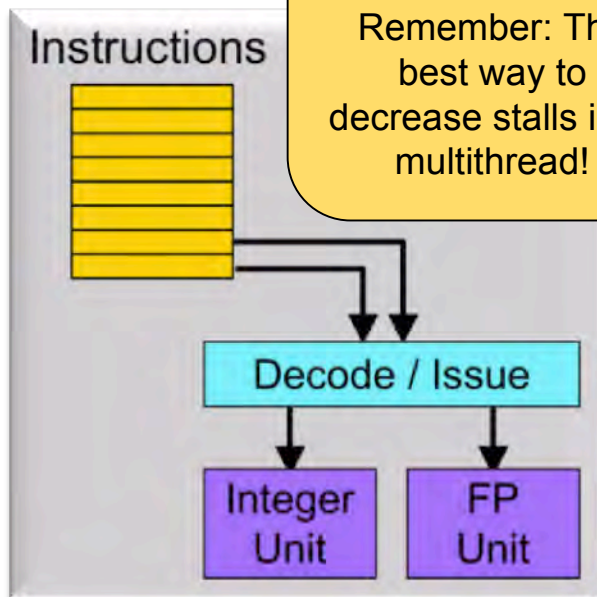
Adding cores alone is inefficient and ineffective.



Each stall can last for 100s of cycles!

Speculative Issue & Multi-Issue?

Remember: The best way to decrease stalls is to multithread!

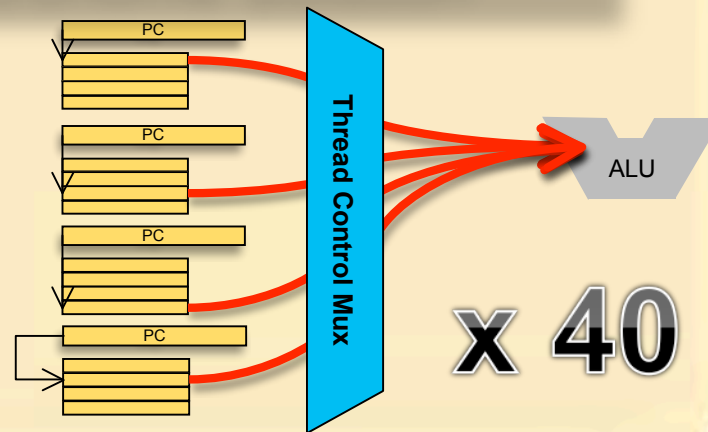


Netronome Solution

✓ Multiple threads per ALU (core)

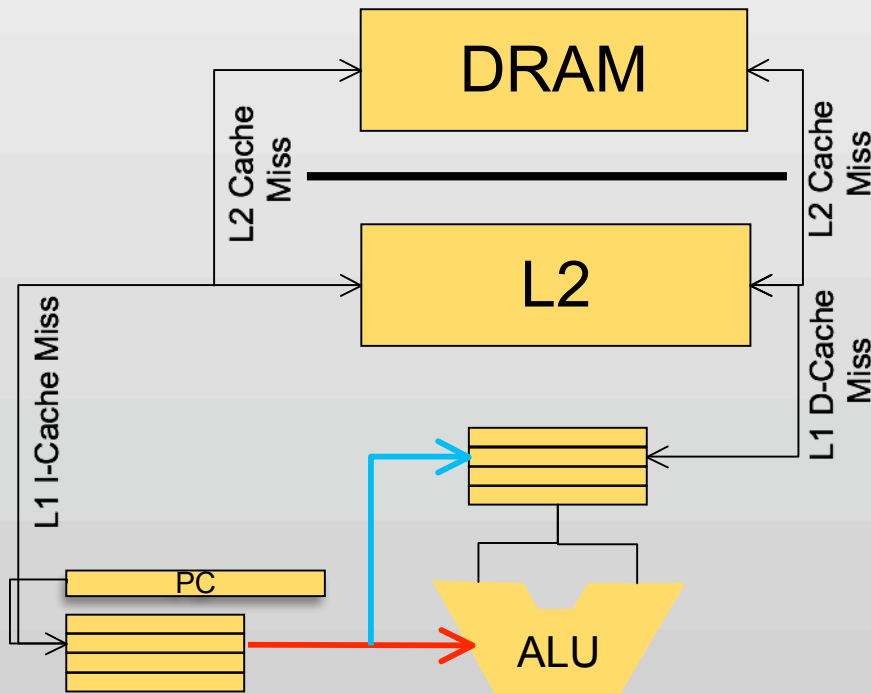
✓ 40 Cores @ 1.4 GHz!

✓ 1884 cycles per packet at 20G!



2: Efficient Memory

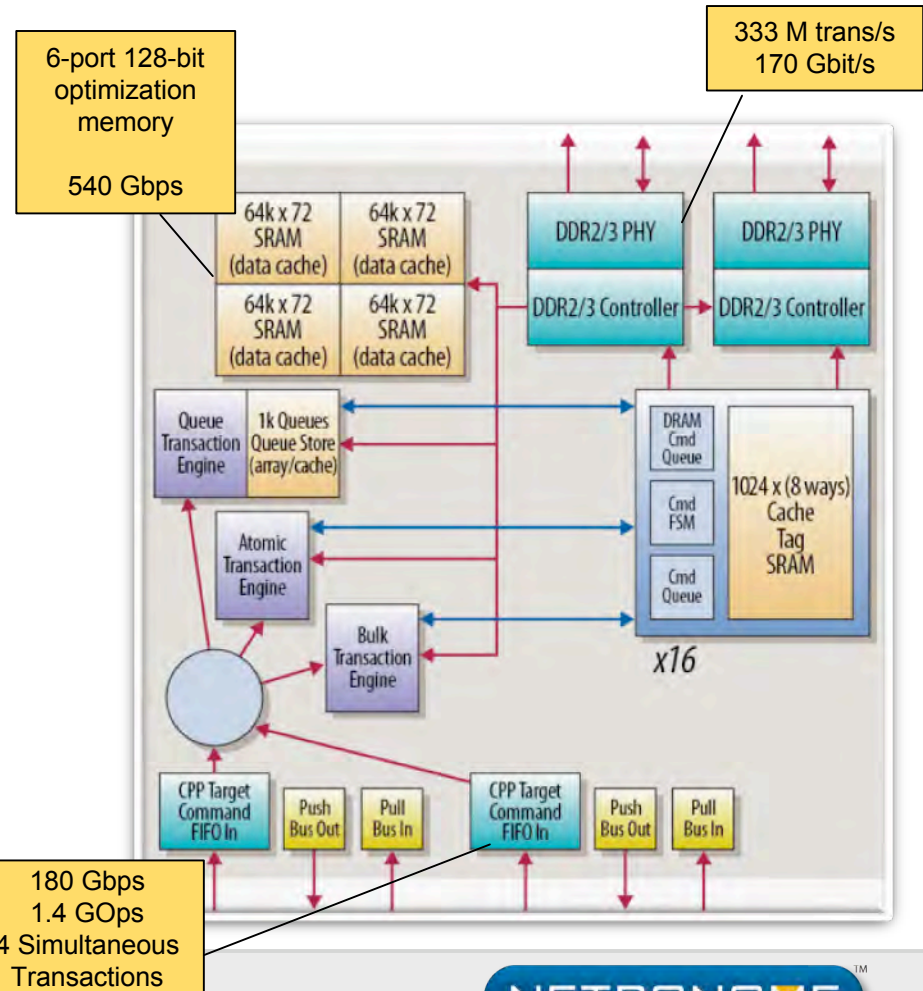
Synchronous Memory Architecture



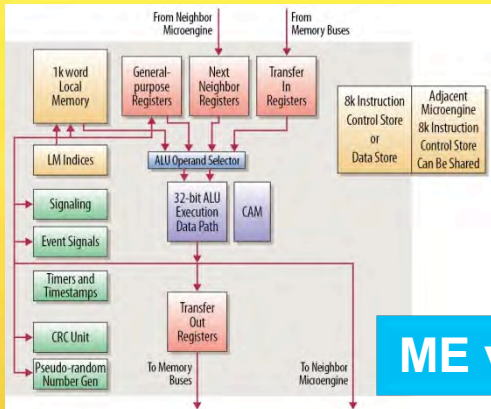
- ALU waits while cache get filled
- Cache fills can be > 100 cycles
- State data overruns small cache
- Result: 10% memory efficiency

Asynchronous Memory

- High-Ratio Multithreading (8 threads / core)
- Internal multi-lane, multi transaction chassis
- Atomic and queue functions
- High-performance external memory
- 96 transactions in flight at one time
- 7x greater efficiency than other methods



3: Core Efficiency



ME v2.7

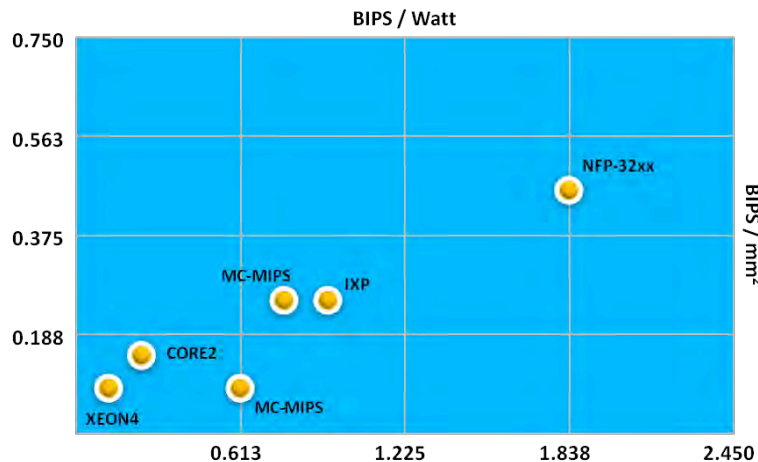
X

40 cores per chip





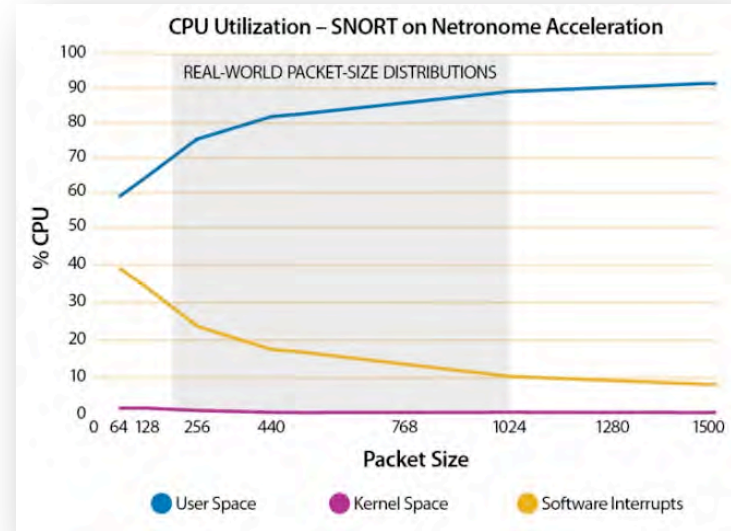

@ 1.4 GHz

- Short 6-stage pipeline
- No bloating speculative execution for line-rate performance
- No bloating multi-issue required for line-rate performance



- Highest performance to power ratio
- Highest performance to silicon space

Case Study: Intrusion Prevention



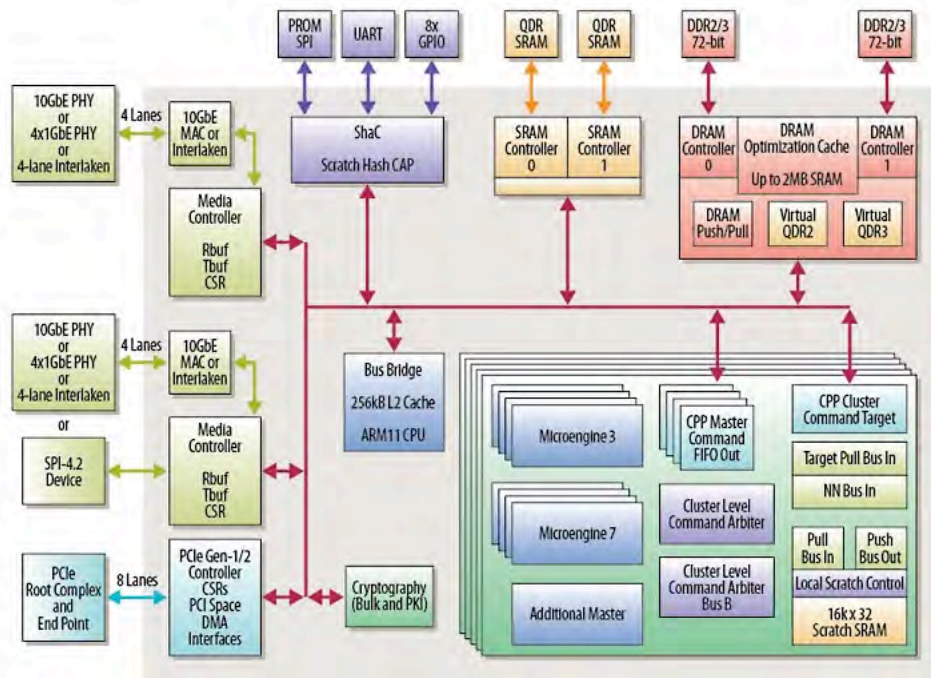
Design Decision Pitfalls

- Look for:
 - Number of cores
 - Ratio of threads to ALU's
 - Number of pipeline stages
 - ALU dependence on cache validity (locality)
- Promises of “easy programming models”
 - Easy to get started, but...hard to get it tuned for real performance
 - Lowest possible instruction thrash
 - High locality of data structures/least thrash possible
- Chip I/O versus what's really inside
 - Remember: 20G I/O != 30 Mpps processing performance
 - Demand full-feature processing at line rate for smallest packets
- Where's the memory
 - Leading multi-core MIPS: 1-2 MB cache / 16 processor cores
 - Leading x86 4-core processors: 8-16 MB of cache

Don't be fooled
by general-purpose architectures
in high-data rate applications!

The NFP is Key to the Heterogeneous Architecture

- Based on Intel's MicroEngine 2.6
12+ years of development
Millions deployed worldwide
- Large software base
- Tight integration with IA/x86 processors



Netronome has taken the most mature networking engine with the greatest number of threads (best memory utilization) and added innovative memory structures and processor offload primitives.

Summary

- The NFP-32xx Network Flow Processor is the industry's first processor purpose built for heterogeneous multicore designs.
- Heterogeneous multicore designs are ideally suited for higher performance dataplane applications
 - Effective Throughput
 - Efficient Memory
 - Core Efficiency
- Multicore x86 provides the most control plane benefits
 - Price
 - Power
 - Performance
 - Functionality
 - Innovation
 - Continuity of Supply

Netronome

- Presenter
 - Kurt Parker
Field Applications Engineering
 - Kurt.parker@Netronome.com
- Full solutions set
 - Network Flow Processors
 - Software Development Kit
 - Software and Libraries
 - Reference Designs
 - Reference Platform

NFP + Multicore x86 > Multicore

Backup Information

NETRONOME PRODUCTS

Netronome Solution

NETWORK FLOW PROCESSORS

The NFP family of network flow processors delivers line-rate throughput, tightly integrated packet and content processing and robust security features in a single chip.



ACCELERATION CARDS

The NFE acceleration cards offer all of the benefits of the NFP flow processors in an easy-to-use PCIe design for quick integration in standard servers.



**Purpose built
for network
coprocessing
in IA/x86
designs**

DEVELOPMENT TOOLS

Netronome's fully integrated development environment (IDE) includes hardware reference designs and standard easy-to-use tools software tools.



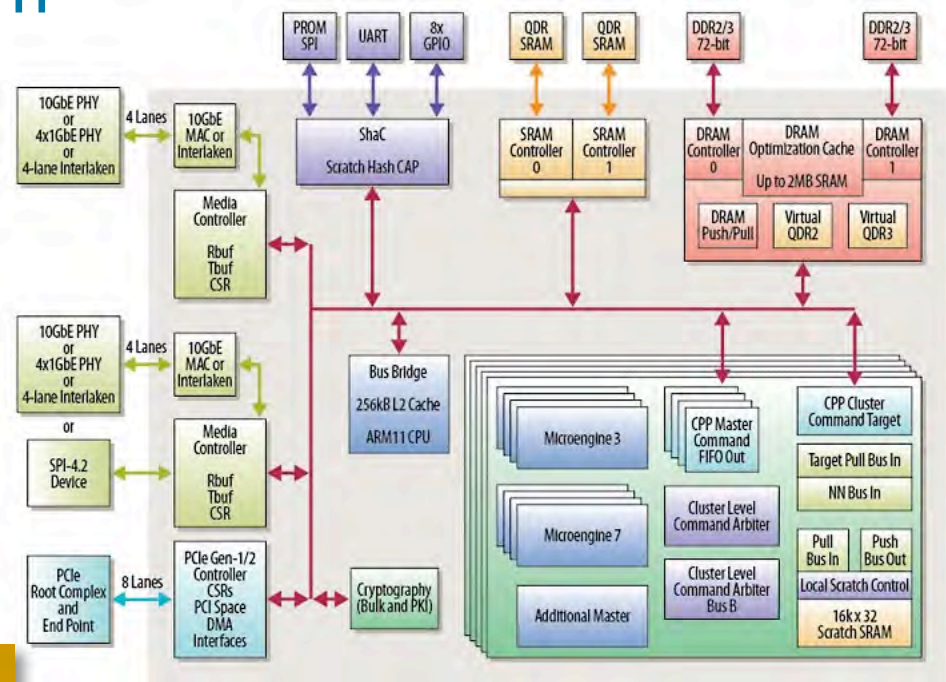
PROFESSIONAL SERVICES

Netronome offers a range of pre- and post-sales design, consulting and engineering support services to ensure that your project is both successful and on-time.



NFP-3200 Block Diagram

40 cores, 1.4 GHz, 56B instructions/sec, 320 threads, 1800 instructions/packet at 30M pps, 20 Gbps of DPI and flow processing.



These enhancements enable Intel® IXP28xx customers to reduce system costs, deliver performance upgrades, extend the life of existing designs and leverage their investment in software, tools and human resources.

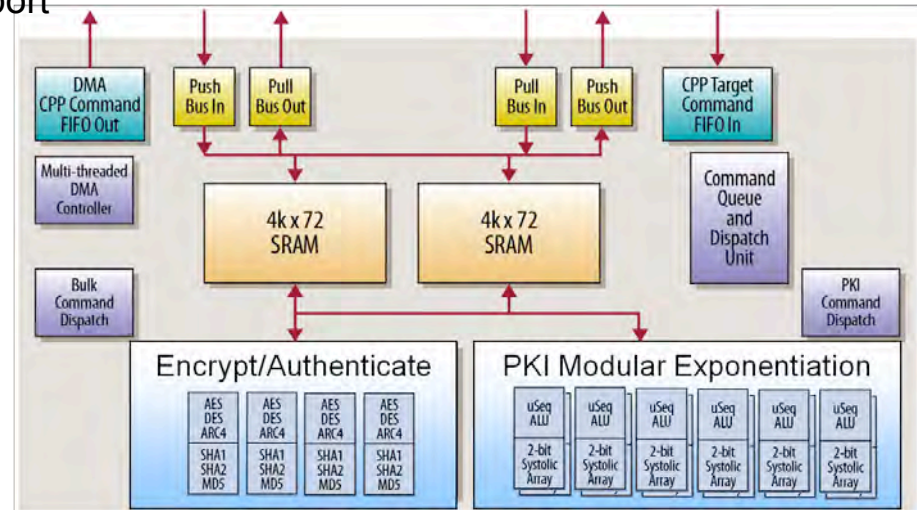
IXP28xx to NFP-3200 Deltas	Customer Benefit:
16 MEv2.5 upgraded to 40 MEv2.7	▶ Almost 3x the processing capability
Replaced RDRAM with DDR2/3	▶ Lowers system cost by more than \$1000
Added Cache to the Memory Controller	▶ Enables higher performance, lower latency
Added Virtual QDR SRAM Interfaces	▶ Enables customer flexibility & lower cost
Increased ME Control Store	▶ Enables more features to be delivered
Replaced PCI with Gen II PCI Express	▶ Enables higher performance
Replaced XScale with ARM11 Core	▶ The L2 cache enables higher performance
Integrated 10Gbps Ethernet MACs	▶ Lowers the system cost and complexity
Added 25Gbps Interlaken to the MSF	▶ Lowers the system cost and complexity
Added high performance flexible crypto hardware acceleration	▶ Provides line rate security acceleration
Integrated PKI H/W Acceleration	▶ Lowers the system cost and complexity

NFP Security Capabilities

- Internal instruction unit
 - DMA, bulk crypt/hash, PKI control, sequenced through cryptography instructions with multithreaded controller
- Hardware accelerated bulk cryptography (20+Gbps)
 - AES-128,192, 256 bit keys
 - ECB, CBC, GCM, CTR, OFB, CFB, CM, f8 support
 - 3DES, DES with
 - ECB, CBC support
 - ARC-4
 - SHA-1, SHA-1 HMAC
 - SHA-2, SHA-2 HMAC 224/256/384/512-bit support
- PKI modular exponentiation
 - 20k+ ops
 - Up to 2048 bit
 - Supports CRT

Many crypto functions for packets destined for software are optionally implemented at packet ingress and egress rather than copy packet into IA memory and invoke a security co-processor which copies it back out and in again

Integrated high performance modern crypto algorithms, with a PKI engine, in a multi-threaded programmable environment

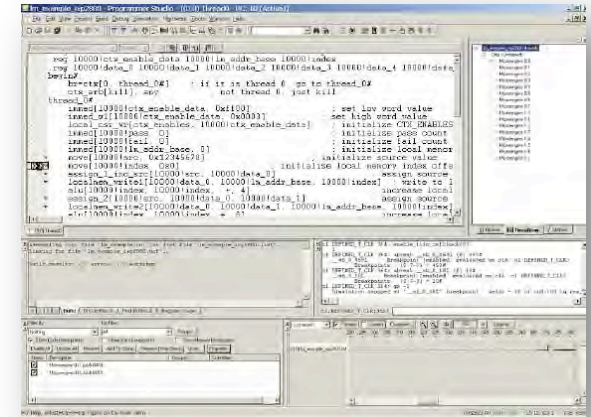


Comprehensive Tools For All Phases of Design



**Integrated
Project Management**

**Project Workspace
with Libraries
and Documentation**



**Centralized Control of Compiler,
Linker, Assembler, Debugger,
Simulation and Testing**



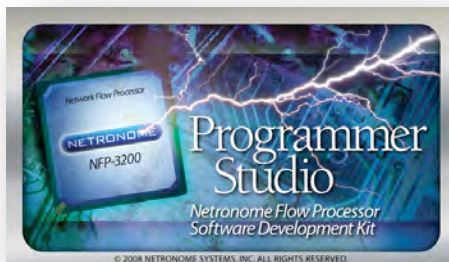
**FlowC Compiler
and Optional
Network Flow Assembler**

**Optional Network Flow
Manager with Sample Code,
Drivers, Libraries and APIs**

**Precision Flow Modeler
for Data and Cycle
Accurate Simulation**

**Local Simulation with Local,
Remote or No Foreign
Model, and Hardware**

**Packet Generation
and Traffic Simulation**

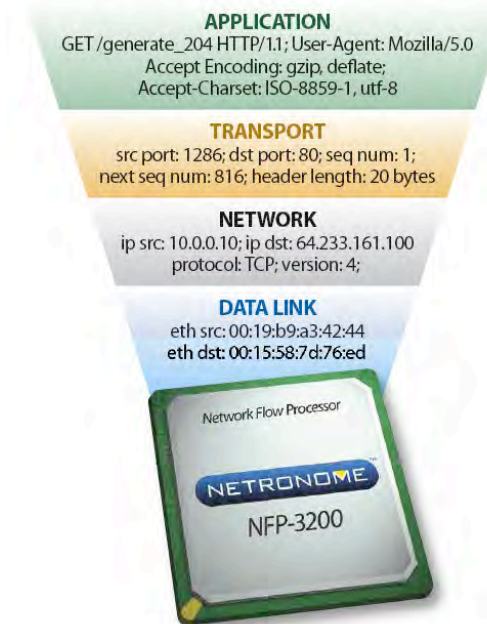


Software: Production and Sample Code

- Production-Ready Distributions
 - Load balancing
 - Packet capture
 - Flow-based DPI
 - Security - IPSec
 - Security – SSL Inspection
 - Available as sample code or APIs
- Neutronome Enhanced Microblocks
 - Packet Rx/Tx
 - L2 switching
 - L3 routing
 - DPI Packet classification, filtering, and cut-through
 - Traffic management and QoS
 - In-line or look-aside security processing
 - Zero-copy VNIC
 - TCP offload
 - Statistics gathering
- Legacy Intel Microblocks



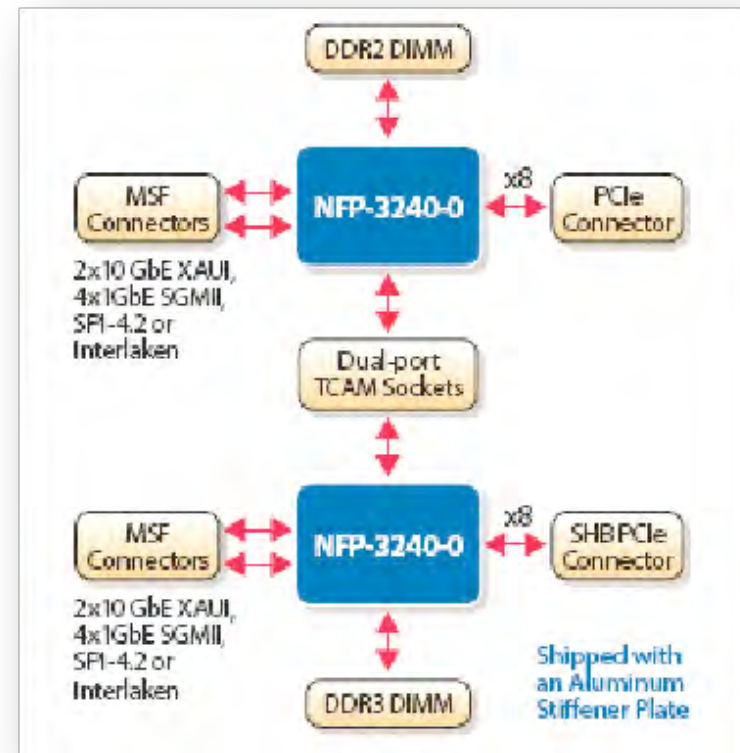
Deep Packet Inspection using NFP-3200



Customer Reference Platform

Base Configuration:

- (1) Dual NFP-3200 Base Board
- (2) 1x10G BaseSR MSF Mezzanine
- (1) MSF Passive Loop Back Card
- (1) PCI-e Passive Loop Back Card
- (1) Dual Port LA-1 TCAM Mezzanine
- (2) DDR2 Memory Module
- (2) DDR3 Memory Module
- System Host Board SHB
- NFP0 to SHB loop card



Downloadable via Support Site:

- Non-commercial Linux Software Board Support Package (BSP)
- Hardware Development Kit (HDK)

Demonstrates NFP performance and functionality while providing customers with a platform and reference collateral to evaluate, prototype and design an NFP-based solution