

# Developing High-performance Embedded Network Security Applications

## *A Heterogeneous Multicore Processing Approach*

The amount of network traffic in today's wired and wireless infrastructure continues to rise at dramatic rates to keep up with demand for IP-based voice/video/data services and applications. It is estimated that annual global IP traffic will increase by 4x growing from 176 exabytes/year to three-quarters of a zettabyte (767 exabytes) by 2014.\*

The major drivers for this incredible growth are primarily video services and mobile data. Video (TV, VoD, Internet Video, and P2P) will exceed 91 percent of global consumer network traffic in four years. Internet-based video will grow from 33 percent of Internet traffic transported to almost 60 percent of all data; taking over as main contributor from P2P traffic (the equivalent of 12 billion DVDs). Mobile data traffic, while still not a large category, will double every year, increasing 39 times during this period.

Network line rates of 10 Gbps are approaching commodity status today. The implication of this incredible growth is that N x 10 Gbps, 40 Gbps and 100 Gbps interfaces will not only be introduced, they will be commonplace in the coming years.

As network throughputs explode, we also need to be able to intelligently monitor our networks for exploits and protect confidential data sources from breaches. An increasing threat to network security is the growing number of high-profile cyber attacks on our communications systems, corporations, financial markets, government agencies and major military installations, and the resultant leakage of classified information and personal data. Even Google has been attacked recently in what could be an instance of state-sponsored corporate espionage. Compounding the problem, there are countless other attacks never publicized, but rather hidden by a web of obscurity for obvious confidentiality and security reasons. The security applications responsible for protecting these critical resources need to keep pace with these increasing network throughputs with even greater network intelligence. Thus, communications equipment must provide complete visibility into network traffic at extremely high bandwidths by using content inspection to ascertain the nature of traffic, not just its destination.

Networks already deploy an array of security applications to protect their critical resources. These applications include virus scanning, firewalls, Intrusion Detection and Prevention Systems (IDS/IPS), Distributed Denial of Service (DDoS) mitigation, Data Loss Prevention (DLP), test and measurement, and network forensics solutions. These applications work almost entirely by providing Deep Packet Inspection (DPI) and flow analysis, looking for known patterns in network flows and blocking or recording them. With the need for application awareness, security processing and DPI, the amount of processing power required for these computationally intense applications grows exponentially at these increasing line rates. These needs

NETRONOME  
WHITE PAPER

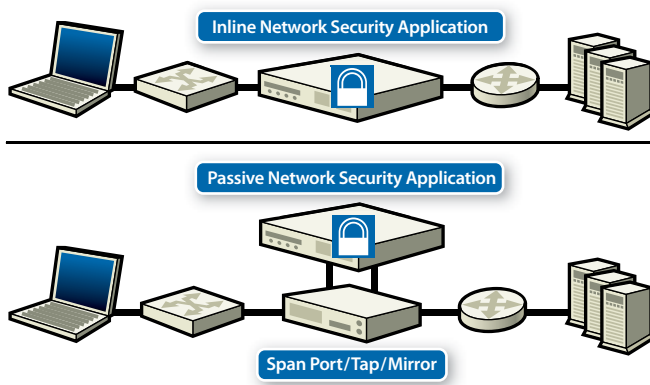
### **EXECUTIVE OVERVIEW**

*Today's network and computing infrastructure is of critical interest to all aspects of our lives, from business and commerce to personal entertainment, and even national security. These vital communication systems share a couple of important characteristics: they require massive amounts of bandwidth to support our insatiable appetite for IP-based services and they require mechanisms offering visibility into all data protocol and application layers to ensure network security. Unfortunately, the network appliances commonly hosting security applications have failed to keep pace with improvements in network performance. However, a new heterogeneous multi-core processing architecture can scale to support tomorrow's throughput needs while providing the ability to see deeply into network traffic.*

NETRONOME

The Flow Processing Company

[netronome.com](http://netronome.com)



**Figure 1: Network and security applications can generally be viewed in several distinct network architectures.**

for increased visibility, throughput and network processing power can be met by a heterogenous multicore processing architecture.

### HETEROGENEOUS MULTICORE PROCESSING PARADIGM SOLVES PARADOX

As shown in Figure 1, network and security applications can generally be viewed in several distinct network architectures. Compute-intensive applications, like intrusion prevention systems, are deployed as active elements sitting directly on the network wire (inline) processing every bit of data that traverses the application in real time. These active security appliances need to operate at network line rates with very low latency. Computation that adds just microseconds of delay to network traffic can ruin the effectiveness of real-time end-user applications like Voice over IP (VoIP) sensitive telemetry systems. Specifically, developers typically view 250 microseconds of delay that any inline network element can add before end-user application performance degrades as a high watermark in 1 Gbps networks.

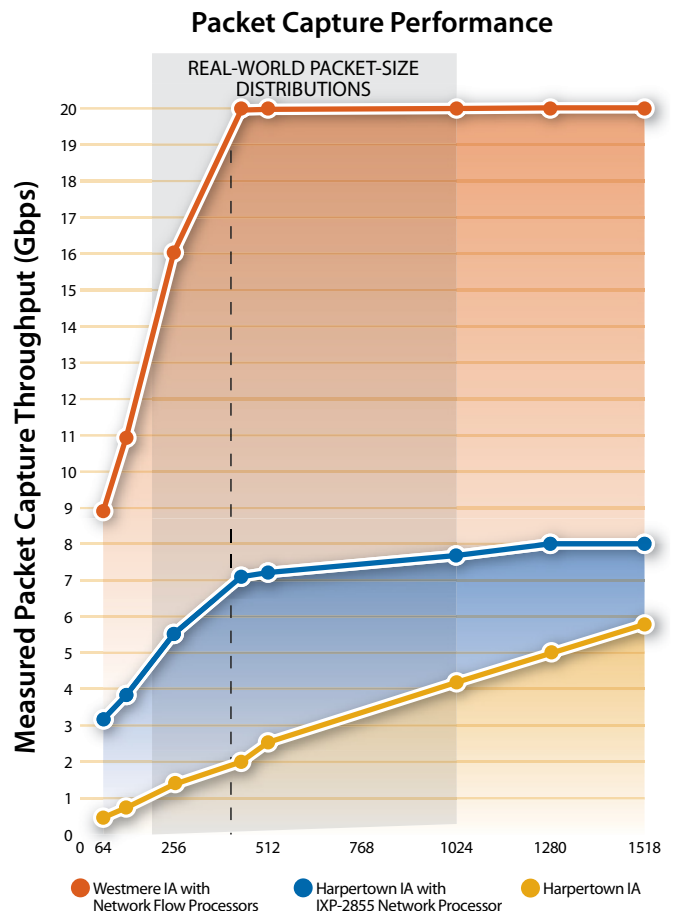
Alternately, passive computing elements like network and computer forensics systems, intrusion detection systems, honeypots and vulnerability scanners are not in the direct network path, but rather are deployed off a span port, network tap or mirrored switch interface. These systems are responsible for collecting and analyzing terabytes of data from distributed sensors deployed throughout the network. These passive monitoring devices can offer a thorough understanding of a network's topology and what services are available, and scan to assess which vulnerabilities might be exposed.

Network appliances deployed in either a passive or active network architecture share a common trait: they must guarantee 100 percent traffic capture across all packet sizes to be effective. Missing any portion of data in a communications stream poses a large threat, making the overriding security application ineffective.

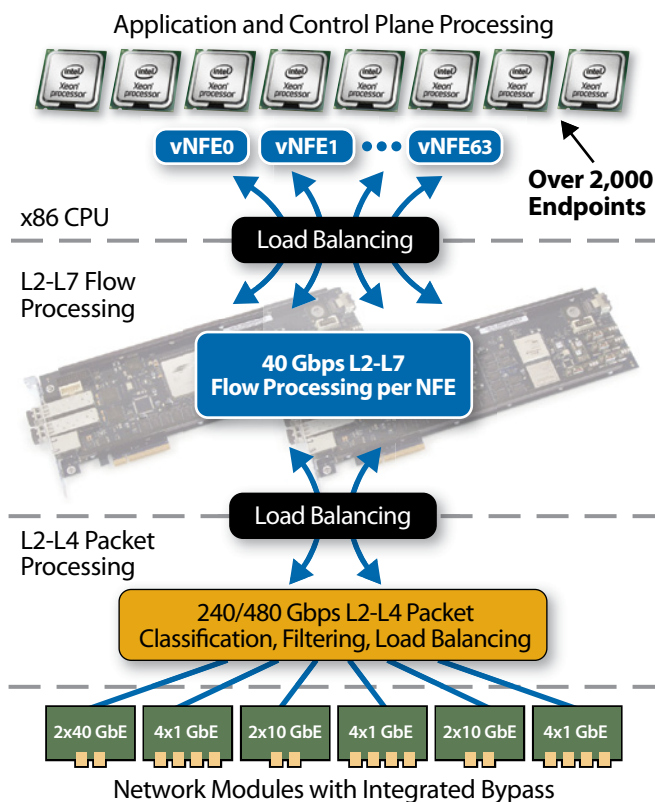
Meeting these performance challenges warrants a new approach to the development of the high-performance systems required by the intelligent network. Such systems need to be capable of analyzing traffic at all

layers of the OSI model, from the data link layer (Layer 2) all the way into the application space (Layer 7) while performing this intelligent processing on all traffic at sustained throughputs of 20 Gbps and higher. Achievement of these goals requires specialized and varied processing elements, each custom-designed for a specific type of workload computation.

A heterogeneous multicore architecture sets a new performance benchmark for embedded application development though separate and discrete processing elements for packet classification, stateful flow management and application and control plane processing, each with increasingly fine granularity. This architecture tightly couples off-the-shelf Ethernet switch processors and network flow processor cores with general-purpose multi-core x86 systems over a high-speed 40 Gbps, virtualized PCIe datapath. This architecture can be scaled from very low-end systems up to appliances offering hundreds of Gigabits per second of packet analysis, stateful flow monitoring, DPI and application throughput, all with a common software architecture. Accelerated designs based on this architecture can enable equipment providers to deliver high-performance, flexible systems that are up to four times more efficient than systems based on x86 general-purpose processors alone with standard Network Interface Cards (NICs), as shown in Figure 2.



**Figure 2: Designs based on a heterogeneous multicore architecture enable high-performance, flexible systems up to four times more efficient than traditional x86 systems and standard NICs.**



**Figure 3: A three-layered heterogeneous processing paradigm uses varied specialized processors to achieve maximum performance while keeping overall system costs low.**

## SPECIALIZED PACKET, FLOW AND APPLICATION WORKLOAD PROCESSING

As shown in Figure 3, a distributed network acceleration architecture uses a multi-chip system to achieve maximum performance and application effectiveness. The three distinct processing stages function as follows:

### Ethernet packet processing

To heighten performance levels, off-the-shelf Ethernet switch processors are commonplace, offering up to hundreds of Gbps of configurable packet processing spanning the datalink, IP and TCP/UDP packet layers. Traffic is classified on ingress and optionally filtered, cut-through to another network interface or load-balanced across the network flow processors (NFPs) that sit logically behind the Ethernet switch processors.

### NFPs to accelerate higher-layer flow processing

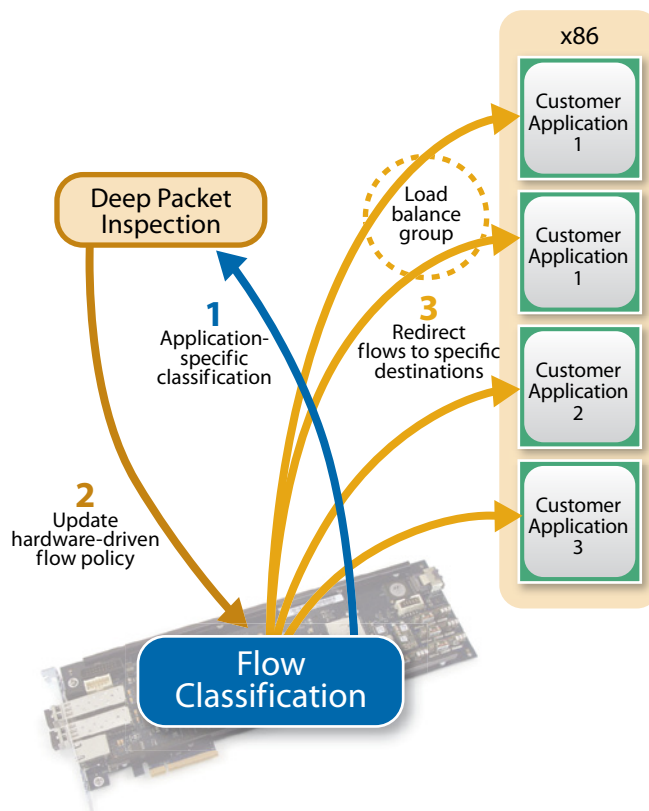
NFPs containing a powerful array of microengine RISC processors are specialized, multicore devices optimized to offload burdensome workloads from general-purpose multicore CPUs. NFPs can handle lower-layer packet processing and accelerate higher-layer flow and application-level processing. This accelerated architecture utilizes the network-optimized NFP cores for switching and routing, packet classification, stateful flow analysis, flow pinning, DPI and dynamic flow-based load balancing. Additional network processing functions, such as checksum offload, TCP termination and SSL processing, can

also be performed on the NFPs and offloaded from the general-purpose CPUs. Traffic can be cleanly structured for transmission from the NFP to the general-purpose cores for application processing, thereby increasing host performance. Additionally, NFPs provide hardware acceleration engines for PKI and bulk cryptography to assure line-rate throughput.

### PCIe communications path to x86 cores

General-purpose multicore x86 CPU(s) in a system are optimized for application and control plane processing. From the NFPs, packets are passed to the x86 cores across a high-performance, virtualization-aware PCIe communications path. An efficient zero-copy technique allows the transfer of packets directly into user-space application memory bypassing the operating system kernel, further accelerating application performance. Flows can be pinned directly to specific applications or load-balanced across parallel application instances to scale application performance.

As shown in Figure 4, through a cooperative set of software APIs between the x86 CPUs and NFP cores, the treatment of flows can also be updated in real-time, offering the ability to change the treatment of a flow after x86 analysis. This type of flexibility is essential in situations where a specific portion of a flow is of interest for inspection. After inspection is complete, all subsequent packets belonging to the flow can be filtered or cut-through at the NFP layer, which conserves valuable



**Figure 4: Via a set of software APIs between the x86 CPUs and NFP cores, the treatment of flows can be updated in real-time. This flexibility is essential when a specific portion of a flow is of interest for inspection.**

PCIe bandwidth and reduces x86 CPU cycles. Through this heterogeneous architecture, the general-purpose multicore processors can focus on the compute workloads they are best suited for, such as behavioral heuristics, Perl Compatible Regular Expression (PCRE) processing, content inspection and analysis or other similar applications.

## INTELLIGENT NETWORKS AT 40 GBPS

The nature of our networks and the important data traversing them creates an opposing set of forces. Networks need to continue to scale to meet exponentially growing bandwidth demands, and enterprises need the ability to effectively monitor these networks at all packet and content layers with stateful network intelligence. To meet these needs, a new distributed, multi-chip, heterogeneous multicore architecture is required, providing specialized workload processing to effectively scale applications to 40 Gbps and beyond.

**Daniel Proch is the director of product management at Netronome Systems, and can be reached at [daniel.proch@netronome.com](mailto:daniel.proch@netronome.com).**



## Accelerate Network and Security Applications to 100 Gbps

*Netronome's network flow processing solution offers the industry's most flexible and highest-performance platform for networking and security applications. With three tiers of packet, flow and application processing, developers can offer over 100 Gbps of throughput in a compact 1U/2U form factor.*

*Netronome's solution combines pluggable 1 GbE, 10 GbE and 40 GbE network modules and Netronome's NFE cards over a high-speed virtualized PCIe datapath with the performance and ubiquity of general-purpose multicore x86 systems.*

\*Source: Cisco Visual Networking Index: Forecast and Methodology, 2009-2014

© Netronome is a registered trademark and The Netronome Logo and "The Flow Processing Company" are trademarks of Netronome Systems, Inc. All other trademarks are the property of their respective owners.  
© 2011 Netronome Systems, Inc. All rights reserved. Specifications are subject to change without notice.  
(4-11)



**Netronome has operations in:**  
**USA** (Pittsburgh [HQ], Santa Clara & Boston),  
**UK** (Cambridge), **Malaysia** (Penang),  
**South Africa** (Centurion) and  
**China** (Shenzhen, Hong Kong)  
**[info@netronome.com](mailto:info@netronome.com) +1 877 638 7629**  
**[netronome.com](http://netronome.com)**