

Networks & Telecom

Packet Processing with Intel[®] Xeon[®]-based ATCA Blades and Servers

ADVANTECH

Enabling an Intelligent Planet

Packet Processing with Intel® Xeon®-based ATCA Blades and Servers


This white paper highlights key design aspects for implementing packet processing applications on Intel® Xeon® processors in AdvancedTCA® and OEM-Ready server topologies.

The dramatic growth currently expected within both fixed and mobile infrastructure, is stimulating a need for new network elements to perform crucial packet processing functions such as security, traffic and packet management. The traffic volumes involved require such platforms to be capable of managing extremely high levels of performance and throughput. The goal of this paper is to provide application and platform developers greater insight as to how packet processing applications can be optimized through the use of Intel® multi-core processors and readily available eco-system software and hardware platforms.

The paper will describe the benefits of using multi-core Intel® architecture processors for packet processing using the latest Xeon® 6-core processors and highlight currently available ATCA node blade designs featuring Intel® Xeon® processors. With respect to networking applications, we will describe and delineate the control and data planes and discuss the role of multi-core processors, the partitioning of cores and the role of the OS. A key focus of the paper will address how through the use of compatible eco-system components developers of packet processing applications can gain significant time to market advantage.

Packet Processing with Intel® Xeon® -based ATCA Blades and Servers

The dramatic growth currently occurring within both fixed and mobile infrastructure is stimulating a need for new network elements to perform crucial packet processing functions such as security, traffic and packet management. The traffic volumes involved require such platforms to be capable of managing extremely high levels of performance and throughput. Much of the growth in mobile networks has been a result of the introduction and dramatic uptake of the new breed of mobile devices such as smartphones, tablets, eBooks etc. According to the Cisco VNI (Visual Networking Index) IP traffic over the next 4 years will grow at a CAGR of approximately 30 percent reaching 22 exabytes per month in 2015.



* Cisco VNI June 2011

To put this into perspective a printed collection of the U. S. Library of Congress would be a mere 10 Terabytes (1 exabyte = 1 million terabytes) and 5 exabytes would represent a transcript of all words ever spoken. This is translating into major investment strategies for carriers and service providers. In August, Dell'Oro Group, one of the key analysts covering the networking and telecommunications industries, reported year

over year revenue increases of 25% in the mobile infrastructure segment. This was driven significantly by equipment destined for mobile data applications. In their overall forecast for the LTE equipment market Dell'Oro anticipates growth at a CAGR of 81% over the next 5 years, and that the total market will reach \$8 billion by 2015. Similar expectations are being seen across all aspects of this network build out especially for all equipment types that will help transport, secure and manage data traffic.

There was a time when all data was simply a stream of bits. Now the huge number of packets traversing our networks each and every day can contain many different types of traffic and the concept of content centric management is extremely attractive, and in many cases necessary. Whether you wish to screen all packets for potential security threats or to choose alternate routing schemes, specialized packet processing platforms are now being built that let you do just that.

Intermingled with all the different types of content packets there are others that contain control information. The volume of control packets is extremely small in comparison to payload data but all packet processing platforms must be capable of making the distinction. Also due to their individual nature they must be handled very differently. Communications networks use a layering concept that identifies and separates these different packets styles as the Control Plane and Data Plane.

Control Plane – From the earliest days of telecommunications the concept of signaling was crucial to the operation of all networks. Using separate networks and protocols of their own (e.g. SS7) these signaling networks provided setup and teardown information for calls, circuits etc. The content of the network control plane is similar in that it provides information and

instructions for a destination platform related to something that must be configured or acted upon. Such packets need to be extracted and sent on to the respective platform's control and management stack. Although the volume of control packets is relatively low there is a more complex processing requirement when compared to the predominantly pass-through data packets.

Data Plane – From original routing architectures this was called the forwarding plane as packets pass through the router on-route to their next hop or destination. Other terms associated with the data plane are fast path, real-time and wire-speed. It is ultimately all the equipment in the data plane that enables (or holds back) the effective speed and bandwidth of a network. One slow network element can hold up everything therefore for data plane related, packet processing equipment, performance is paramount. Hence the term wire speed, describing the real time nature of the traffic flow and the expectation that any processing and content related decision-making all happens without impeding traffic flow in the slightest way.

The data plane is where the money is. The traffic transported across the 100s of millions of network nodes that make up the Internet (and private networks) has much diversity. Traffic types will have varying levels of criticality and importance to their respective users. Some connections are delivered using 'best efforts' others are subject to specific QoS (Quality of Service) and QoE (Quality of Experience) metrics and may even be subject to stringent SLAs (Service Level Agreements). For example, three sequential packets in a data stream may contain part of an email, a web page and a phone call. A five second delay in the delivery of an email is unnoticeable, annoying in the case of a web page but would cause unacceptable quality issues in a VoIP call. What if one were able to identify these different packets and potentially route each to infrastructure tuned for specific performance criteria. Now consider being able to identify specific traffic types and take further actions based on what you may find e.g. for security

reasons such as virus/intrusion prevention or legal intercept. This is all now possible through the implementation of applications that utilize Deep Packet Inspection technologies.

Packet inspection and filtering techniques have been utilized for a while but these have primarily looked at packet headers and filtered based on port classifications etc. Deep Packet Inspection or DPI is a technique that will literally look inside each and every packet in a data stream. The content is then compared to known example libraries and once identified the application can take whatever action it deems appropriate.



Through the application of new DPI technologies a major increase in the granularity of traffic management and control is now possible. For example, carriers can now improve revenue by being able to offer different levels of service on a per user basis, dependent on traffic type and even destination or time of day. New 4G mobile infrastructure has Policy and Charging Control (PCC) functions that provide operators with advanced tools for service-aware Quality-of-Service (QoS) and charging control. Within the network, the Policy and Charging Enforcement Function (PCEF) identifies and associates applications and/or users with specific traffic flows, applying policies to individual sessions based on requirements defined by a network element called the Policy and Charging Rules Function (PCRF). All of this is possible through the

use of DPI implemented on high performance packet processing platforms.

In theory, executing a series of content lookups doesn't sound too hard. While it is somewhat more complex than it sounds the challenges relate more to throughput efficiency than process complexity. Thus we can say that DPI represents a huge volume of relatively repetitive 'simple' tasks that must execute very quickly, in fact, at better than wire speed. This all makes the efficiency of such packet processing platforms the most critical factor. Coming back to the data and control planes we must now separate them not just logically but architecturally into the fast and slow paths.

The controlling entity of any compute platform is the operating system (OS). While all-encompassing and capable of managing everything from user code execution to device control and the management of storage and network I/O, operating systems are inherently inefficient – the slow path. Even with complex schedulers and interrupt management schemes there are significant unnecessary overheads that make it undesirable to attempt routing fast path traffic through the OS. Earlier packet processing implementations were built using single core processors. With only one processor in use there was no way to avoid the slow path without adding more processors whether general purpose CPUs, specialized network processors or ASICs. With the advent of multicore processing technologies the available options were increased. Multicore architectures enable processors to be created that have two or more identical CPU cores and typically share a common system memory. Each core can operate independently on different processing elements and data flows, and can also easily interact with

other cores and processors. As we mentioned earlier, DPI applications benefit when the fast path processes can be separated, they are also easily split into logical chunks with the heavy lifting processes being highly repetitive making it ideal for scaling across many cores.

Multi-Core Processors

Intel® has a broad range of high performance multi-core processors. For embedded computing and communications applications, such as DPI, there are 8 variants as part of the Intel® Xeon® processor 5600/5500 series. There are options with 2, 4 or 6 cores and thermal design power (TDP) ranging from 38W to 80W. Four processor options provide robust thermal profiles, ideal for the AdvancedTCA form factor and applications requiring compliance with NEBS Level 3 thermal specifications.



Multi-core designs don't in themselves enable greater performance unless you can feed and manage them efficiently and optimize software execution to gain every last ounce of benefit from those extra cores. The Intel® Xeon® 5600/5500 series processors have a number of innovative enhancements that deliver on the promise of multi-core execution.

- Greater Parallelism increases the amount of instructions that can be run "out of order." This enables more simultaneous processing and overlap latency. To be able to identify more independent operations that can be run in parallel, Intel® increased the size of the out-of-order window and scheduler, giving them a wider window from which to look for these operations. Intel® also increased the size of the other buffers in the core to ensure they wouldn't become a limiting factor.

- Intel® Hyper-Threading Technology enables more energy efficient means of increasing performance for multi-threaded workloads. This next generation microarchitecture's SMT capability enables running two simultaneous threads per core.
- Application Targeted Accelerators extend the capabilities of Intel® architecture by adding performance-optimized, low-latency, lower power fixed-function accelerators to benefit specific applications.

Silicon solutions lie at the heart of all packet processing platforms. There is a complete integrated ecosystem that builds upon the silicon to provide the overall platform solution. Starting with the silicon the value chain grows to include all these key players; Silicon, Blade or Platform Hardware, Integration Services, Operating System, Network Stacks and Protocols, Enabling Middleware, Application Solution Integrators and VARs. Let's take a look at some of these crucial pieces of the puzzle.

Intel® Data Plane Development Kit

A number of these value chain elements combine to create the software infrastructure on which DPI technology based solutions are created. In order to help accelerate the whole ecosystem, Intel® has created a development kit with support for their processors such as the 5600/5500 series. The Intel® Data Plane Development Kit (DPDK) has been ported and integrated with specific offerings from leading Intel® ecosystem partners who have worked together to make this available on their respective hardware and software platforms. The DPDK is a collection of Data Plane libraries that cover the areas of queue and buffer management, packet flow classification and poll mode drivers. Packaged with optimized NIC drivers and built on an environmental abstraction layer

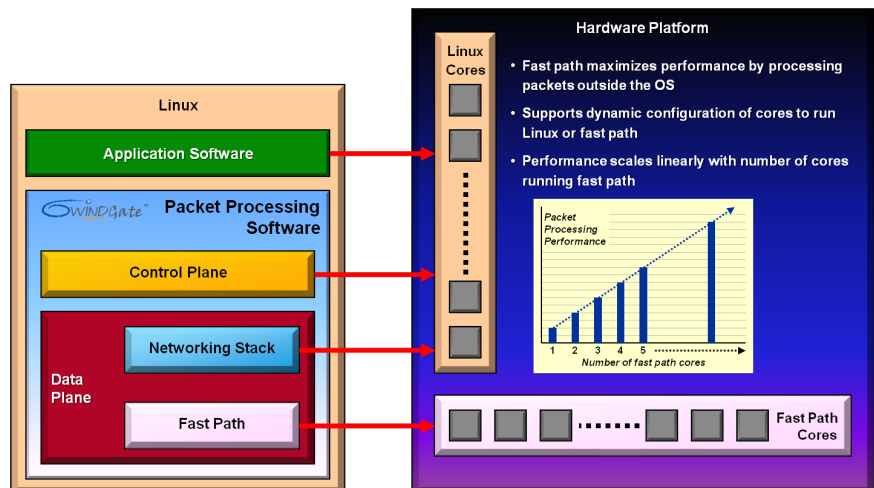
that allows for easy porting and integration with supported IA platforms the DPDK creates a solid foundation and starting point for developing packet processing solutions.

One further step up the value chain are fully featured products that have been ported to the Intel® DPDK. One such product is 6WIND's 6WINDGate that replaces standard operating system networking stacks and using multi-core architectures, deals with the inherent OS inefficiencies we discussed earlier.

6WINDGate Packet Processing Software

The 6WINDGate™ packet processing software product implements networking stacks that are split into two layers. The lower layer, typically called the fast path, processes the majority of incoming packets outside the OS environment and without incurring any of the OS overheads that degrade overall performance. Only those rare packets that require complex processing are forwarded to the OS networking stack, which performs the necessary management, signaling and control functions. In a typical networking application, when the standard OS networking stack is replaced by optimized packet processing software based on a fast path concept, the networking performance of the processor subsystem will increase by seven to ten times.

With a comprehensive set of protocols available



for the control plane, the networking stack and the fast path, 6WINDGate provides developers with a single-vendor solution for all the protocols required for a high-performance networking platform based on multicore technology. 6WINDGate is a drop-in replacement for standard Linux networking stacks and is fully-compatible with Linux application APIs. Any application software that has been developed to use Linux networking APIs will run unmodified when 6WINDGate is used.

When installed on a multicore platform, 6WINDGate can be configured at run-time to make the optimum use of the number of cores available. For example, on a six-core Intel® Xeon® processor E5645, the system could typically be configured as follows:

- One core is configured to run Linux, as well as the 6WINDGate control plane and the 6WINDGate networking stack
- The remaining five cores are configured to run the 6WINDGate fast path, which makes full use of processor-specific services provided by the Intel® Data Plane Development Kit (Intel® DPDK) software.

The allocation of logical cores either to the control plane and the networking stack or to the fast path can be changed dynamically as traffic patterns evolve. Using IPsec as an example, the control plane requires a significant amount of processing capabilities while all IKE sessions are being established. Once the establishment process is completed, all the cores except one can be used by the fast path for IPsec processing for the traffic. This avoids over-provisioning some CPU resources for the control plane and for the fast path.

Qosmos Network Intelligence Technology

With the fast path appropriately optimized the next step is to add a comprehensive DPI development environment. Qosmos specializes in unique network intelligence technology that provides unprecedented visibility into data traffic. The company delivers software development kits and intelligent IP probes which recognize thousands of protocols and metadata attributes for the most accurate picture of real-time data activity on networks.

Qosmos ixEngine Supported Protocol Families and Metadata

- **Networking:** Ethernet (20 metadata), IPv6 (23 metadata), UDP(21 metadata), TCP (35 metadata), DHCP (66 metadata) and over 70 other networking protocols
- **Email:** IMAP (71 metadata), POP3 (52 metadata), SMTP (66 metadata), and 5 other email protocols
- **Webmail:** Yahoo!Mail (64 metadata), Gmail (64 metadata), Windows Live Hotmail (54 metadata), and more than 20 other webmails
- **Instant Messaging & Multimedia Conferencing:** Skype (15 metadata), Windows Live Messenger (51 metadata), Gtalk, AIM and 20 other IM protocols
- **Forum:** Google Groups (46 metadata), Yahoo Groups (37 metadata), MSN Groups, and several other forum applications
- **VoIP:** SIP (66 metadata), RTP (27 metadata), H225 (27 metadata), and 8 other VoIP protocols
- **Web browsing:** HTTP (69 metadata including http request, URL, DNS, cookies, referer, Device and browser), HTTPS
- **Social networks:** Facebook (36 metadata), Twitter, LinkedIn, MySpace and 155 other social networks
- **Web multimedia streaming:** Flash /RTMP (15 metadata), RTSP (used by Real Player and QuickTime, 60 metadata) , MPEGTS, Youtube, Dailymotion, and 10 other streaming protocols
- **Search engines:** Google (20 metadata), and 5 other search engines
- **E-commerce: eBay (17 metadata), and 5 other sites**
- **Websites:** Google Maps (23 metadata), Megaupload, and 50 other websites
- **Mobile networks:** GTP (57 metadata), MMSE, and 8 other protocols used in mobile networks
- **Peer to peer:** Bittorrent (43 metadata), eDonkey, and 15 other P2P protocols
- **File management:** FTP (24 metadata), NFS, and 13 other protocols
- **Enterprise:** MySQL (16 metadata), CUPS, McAfee and 30 other protocols
- **Gaming:** World of Warcraft (13 metadata), Counter Strike
- **Authentication:** Radius (30 metadata), Socks4

Qosmos' ixEngine is a DPI Software Development Kit (SDK) composed of software libraries and tools that are easily integrated into new or existing solutions. ixEngine allows applications to be built that can perform real-time packet classification at protocol and application levels and extract metadata. The software has many advanced capabilities including: parsing of tunneling protocols, data filtering, and visibility at ALL levels (flow, session, IP application, and subscriber). At the core of any DPI functionality is the ability to identify and isolate specific traffic types. ixEngine's libraries provide instant access to 1000+ protocols and 4500+ communications metadata.

Technology evolves and so do protocols, threats and traffic types. Qosmos has an automated and proactive update service. Through their Protocol and Application Watch service they are able to continuously add new protocols and update existing protocols and applications in the Protocol Plugin Library.


Although the Qosmos Protocol Plugin Library covers hundreds of protocol and application signatures there is always the possibility of custom or regional variants. Recognizing this, Qosmos offers a service for on-demand development of protocol and application signatures. There are also situations when OEM developers may prefer the flexibility and independence to develop their own custom protocol plugins. Qosmos developed their Protocol Plugin Creator that uses the Qosmos Protocol Description Language (PDL) to provide just such a capability.

For Networking OEMs that are building out their packet processing capabilities, utilizing field proven software/middleware such as the offerings from 6WIND and Qosmos provides an important time to market advantage. The OEM can focus on their specific packet processing application relying on the underlying APIs and software functions to provide the necessary DPI capabilities and packet throughput optimization. All of these software elements must have a

hardware platform on which to run, first for development and then for deployment.

Advantech High Performance Blades and Appliances

Packet processing applications using DPI technologies come in many shapes and sizes from security and threat management appliance style platforms to large, rack based, traffic management solutions for 4G mobile networks. The Intel® Xeon® processor 5600/5500 series have specific versions that were designed to support the needs of embedded communications applications including thermal profiles that make them ideally suited for AdvancedTCA form factor blades. Supported as part of the SCOPE Alliance's profiles and carrier grade base platform definition, numerous network platforms have been built using AdvancedTCA. Advantech is a Premier member of the Intel® Embedded Alliance and the Networks and Telecom group has a broad selection of platform and blade products that support these Intel® embedded multi-core processors and the software environments detailed earlier.



MIC-5322 Dual Intel® Xeon® 5600 Series Processor Blade

The MIC-5322 is a dual processor Intel® Xeon® 5500/5600-based ATCA blade. It enables the highest performance available in ATCA form factor, with 12-cores and 24-threads of processing power, low DDR3 memory latency, fast PCI Express 2.0 and accelerated virtualization. The Intel® 82599 10 GbE controller

plays a key role in end-to-end network performance and throughput, including a PCIe x8 interface (40Gbps) to improve the entire data path as well as multi-core optimized queue support. The flexibility of the Intel® Xeon® 5500/5600 Series allows tremendous upgradeability, scalability and cost efficiency options with two, four or six-core processors fully supported.


- Two 2, 4 or 6-Core Intel® Xeon® 5500 or 5600 processors
- Intel® 5520 IOH36D/ICH10R server class chipset
- 6 DDR3 VLP DIMMs up to 48 GB with ECC support
- Two XAUI ports on Fabric interface
- Two 1000BASE-T ports on Base interface
- Two 1000BASE-T front panel ports
- Two USB2.0 front panel ports
- Fully managed, hot swappable RTM

The dual processor MIC-5322 and its single processor companion the MIC-5320 share similar specifications, both offering the highest available AdvancedTCA performance envelope and network throughput fueled by a 10GbE controller. For complex lookup applications such as DPI the highest levels of performance is assured through the 48GB of triple channel DDR3 making Advantech's Intel® Xeon® Processor 5600-based ATCA blades the optimal choice for large scale packet processing platforms.

Dual Intel® Xeon® Processor 5600 Series Based 2U Network Appliance

For smaller scale, appliance style, implementations or as a software development platform, Advantech offers the FWA-6500, an Intel® Xeon® Processor 5500/5600 Series 2U 10 GbE Network Server. Built with flexibility and scalability in mind, the FWA-6500 system provides the high performance and the feature set required for dual processor-based high-end network appliances with configuration options optimized for intensive computing, energy-efficiency and high speed I/O. The Intel® Xeon® processors are based on the Intel® QuickPath

Interconnect architecture and VT technology improves the performance in a virtual environment.



The FWA-6500 supports DDR3 1066/1333 Registered ECC memory with the memory controller embedded in the processor. The chipset consists of the Intel® 5520 IO Hub (IOH), Intel® I/O Controller Hub (ICH10R) and the I/O subsystem. The Intel® 5520 IOH provides 36 PCIe Gen2 lanes and the ICH10R provides 6 x PCIe Gen1 lanes. Therefore, the FWA-6500 can provide different combinations of Network Mezzanine Cards such as 4 x GbE (RJ45/SFP), 2 x 10 GbE (SFP+) and 8 x GbE (RJ45/SFP). In addition, the FWA-6500 can support standard PCI-express expansion slots with a riser card.

- Two 2, 4 or 6-Core Intel® Xeon® Processor 5500 or 5600 series
- 12 x DDR3 1066/1333 Registered ECC Memory up to 96 GB
- 4 x Network Mezzanine Cards (NMC) using PCIe x8 gen.2 connectors
- NMC FRU modules include 4-port GbE (RJ45/SFP), 2-port 10 GbE (SFP+) and 8-port GbE (RJ45/SFP)
- 2 x PCIe full-height/ half-length add-on cards
- 2 x 2.5" removable SATA HDD (3.5" removable SATA HDD option)
- Supports Remote Management IPMI 2.0

In Conclusion

The technologies we have seen throughout this paper represent the leading edge in silicon, software and platform level designs. New capabilities in areas such as Deep Packet Inspection technology will be crucial to the evolution of the next generation of high performance networks. Companies such as Advantech, Intel®, Qosmos and 6WIND have created solutions that will work together to create advanced packet processing platforms. These platforms will then act as the foundation for service provider and carrier network elements to address the challenges and opportunities that have risen from the ongoing growth experienced throughout global data networks.

Advantech Contact Information

Hotline Europe: 00-800-248-080 | Hotline USA: 1-800-866-6008

Email: NCG@advantech.com.tw

Regional phone numbers can be found on our website at <http://www.advantech.com/contact/>

www.advantech.com/nc

