

# Adopting Multicore Processors for 40G Telecom Equipment

Charlie Ashton, 6WIND and Gene Juknevicius, GE Intelligent Platforms

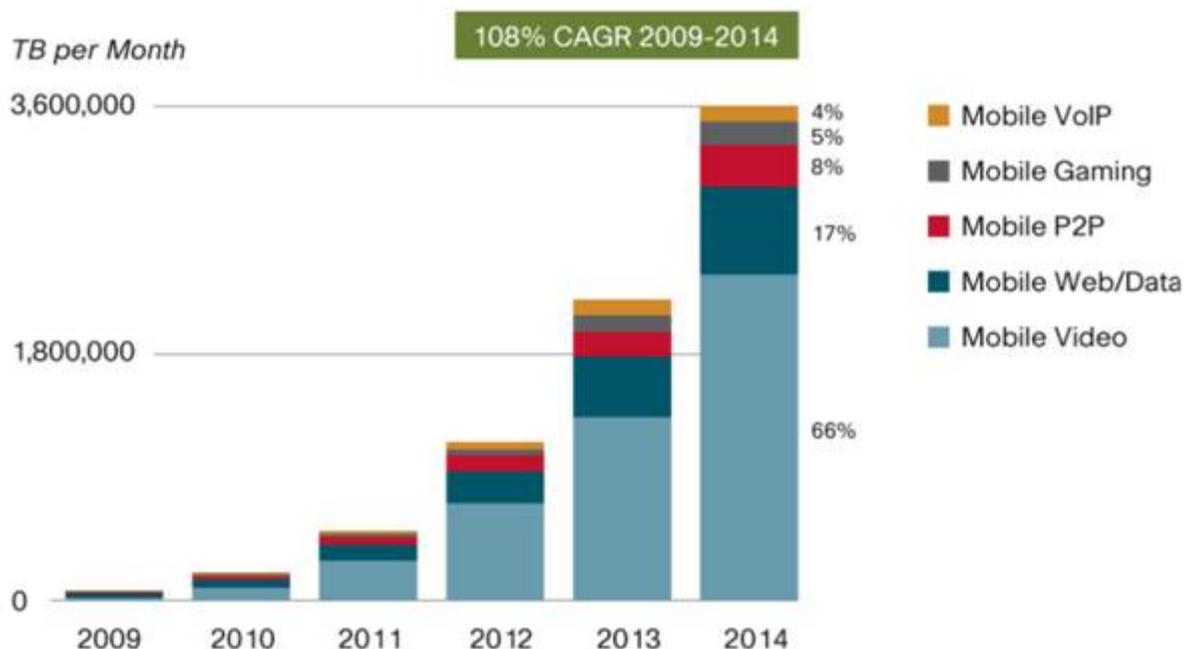


Driven by the needs of high-bandwidth Internet applications, next-generation telecom blades are required to achieve 40Gbit/s throughput for Ethernet and IP-based traffic. While several suppliers provide multicore processors optimized for these high-performance systems, each is based on a proprietary architecture and there is no software compatibility between processors from different suppliers. Furthermore, in order to properly evaluate performance a test application needs to take advantage of hardware-specific acceleration features, which increases engineering effort prior to final hardware selection. As a result, system designers are at risk of having to select a processor architecture early in their development cycle, with limited scope for making changes later if that processor is delayed or fails to meet initial feature/performance goals.

GE's ATCA blades, with ready-to-use software from 6WIND, enable customers to quickly set up and run performance tests using a number of high-level protocols. Since the 6WINDGate software takes full advantage of hardware-specific acceleration features, customers' initial test results will reflect high performance with the accuracy that is needed for realistic applications. In addition, 6WINDGate provides an effective hardware abstraction layer, which allows customers to painlessly migrate their application from one GE platform to another, mitigating the risks associated with being locked into a particular processor family.

# Adopting Multicore Processors for 40G Telecom Equipment

Published on Electronic Component News (<http://www.ecnmag.com>)



Source: Cisco VNI Mobile, 2010

Only a few years ago, 10Gbit/s Ethernet was the state-of-the-art in ATCA systems for telecommunications. Today, however, 40Gbit/s systems are in development with 100Gbit/s data rates on the horizon. At the same time, the level of required processing continues to increase. Security and privacy concerns require data to be encrypted and continuously searched for possible virus signatures. A number of protocols are encapsulated together, requiring efficient packet parsing and header manipulation capabilities. On the hardware side, a number of suppliers such as Cavium, Intel and NetLogic offer processors that are positioned for packet processing at 40Gbit/s data rates. In order to increase performance within these processors, certain tasks are offloaded to hardware acceleration engines, such as encryption/decryption, regular expression search, and data compression/decompression. Such hardware offload engines deliver very high performance, although their usage requires software to be specifically written to take advantage of them.

Another technique used by processor suppliers in order to deliver higher packet processing performance is providing feature-reduced and performance-optimized operating systems, so called "bare metal" operating systems or Multicore Execution Environments (MCEEs). For example, Cavium developed their Simple Executive OS and Intel provides their DPDK® software. Applications written for bare metal operating systems run on individual processor cores and are highly efficient for packet processing tasks. Since bare metal operating systems do not use context switching or interrupts, applications typically run to completion and deliver very predictable and repeatable packet processing performance and latency. For instance, a simple packet forwarding function written in the Cavium Simple Executive can perform packet forwarding with latency under 2µs, and this value is highly repeatable.

Although bare metal operating systems support programming in the standard C

# Adopting Multicore Processors for 40G Telecom Equipment

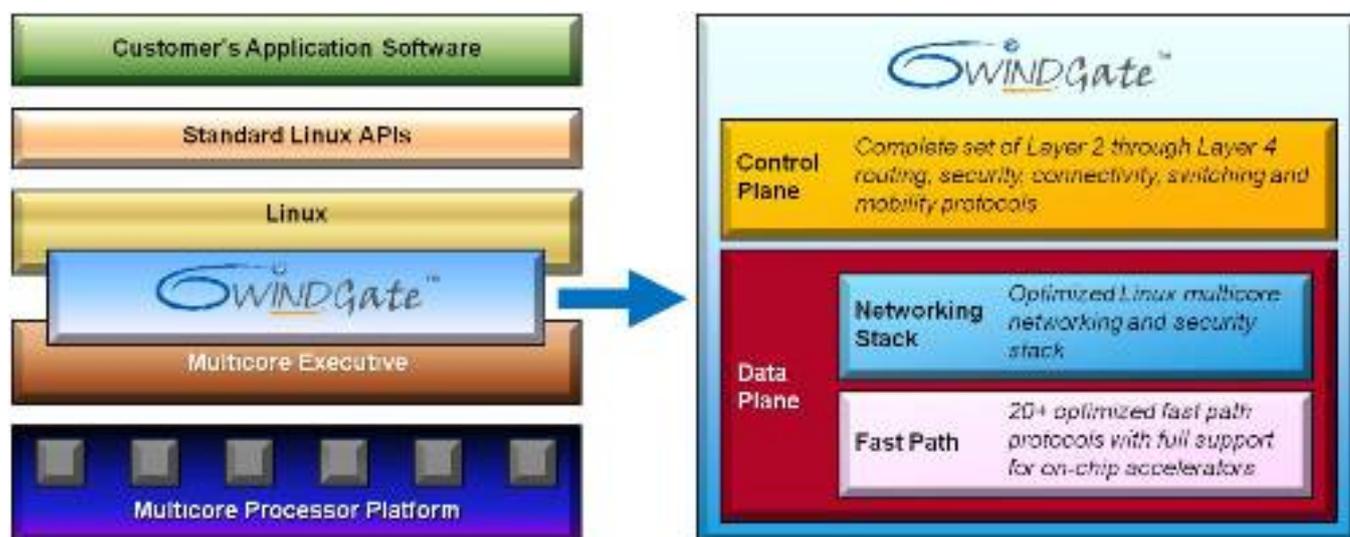
Published on Electronic Component News (<http://www.ecnmag.com>)

language and include some basic software libraries, they lack an extensive software library and the protocol support that is available for Linux. Therefore, implementing sophisticated packet processing directly in a bare metal operating system requires significant programming effort and expertise in processor-specific features.

Furthermore, applications written in such a way will be tightly coupled to a specific processor, making migration to different architectures painful and time-consuming. This effectively locks telecom equipment suppliers (TEMs) into a specific processor family, leaving them at risk of changes to their processor supplier's new product roadmap or delivery schedules. The schedules for releasing complex processors often slip and the several month delays commonly seen in rolling out new processors can significantly affect a TEM's new product rollout schedule and rhythm.

To mitigate these risks, customers are looking for ways to reduce dependency on a particular silicon architecture and to reduce their software programming efforts, while still maintaining high performance. Such a goal is difficult, but can be achieved by introducing a hardware abstraction layer. Ideally this would provide a Linux-like application development environment that supports a large number of common Ethernet and IP protocols, while delivering the packet processing performance that can only be achieved by using a bare metal operating system and leveraging processor-specific hardware offload engines. Furthermore, such a hardware abstraction layer would reduce the customer's dependency on the underlying processor architecture, allowing hardware to be changed without a major application software redesign.

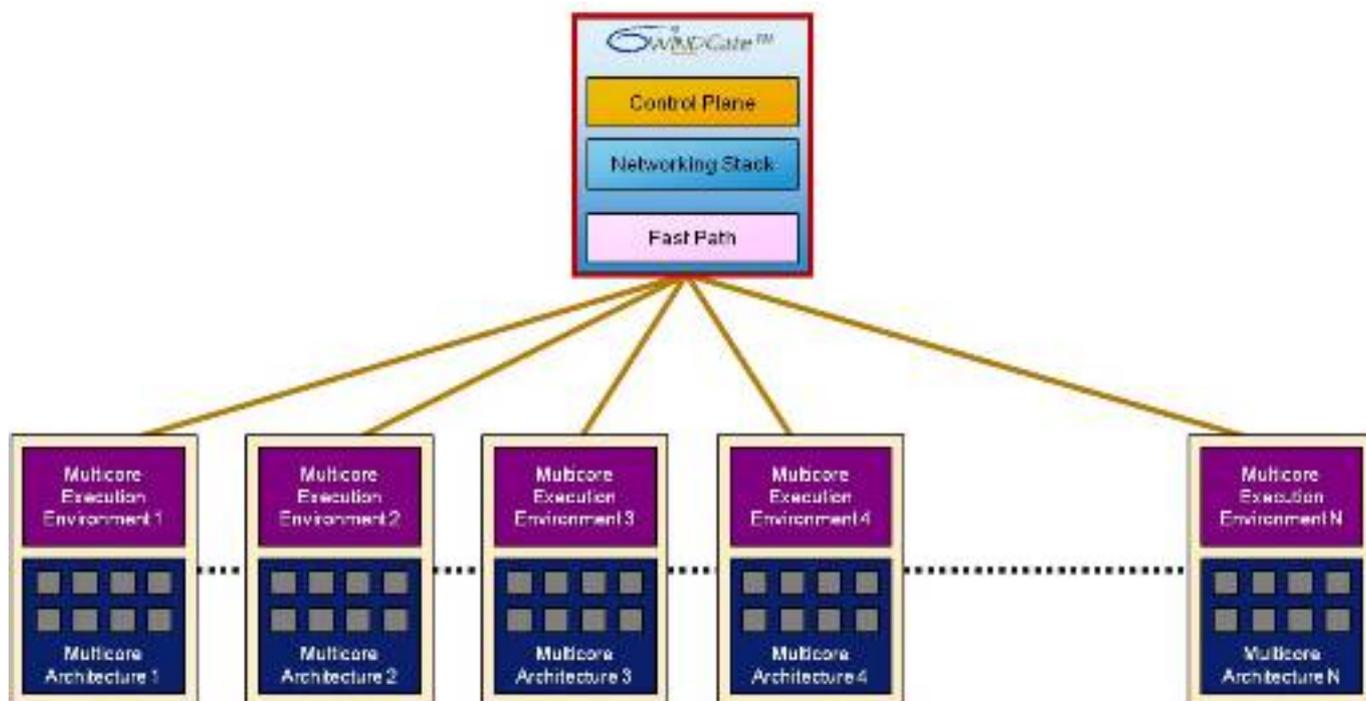
As illustrated in the diagram below, the 6WINDGate software is a drop-in replacement for the standard Linux networking stacks and is fully compatible with standard Linux APIs, regardless of which multicore processor target is selected. Any application software that is developed to use standard Linux networking APIs (Netlink, PF\_KEY, Netfilter, BPF/tcpdump etc) will continue to run unmodified when 6WINDGate is used in the system. This enables TEMs to maintain a single, unified code base in the confidence that it will run correctly on whatever 6WINDGate-compatible processor platform they select for a specific end-product.



# Adopting Multicore Processors for 40G Telecom Equipment

Published on Electronic Component News (<http://www.ecnmag.com>)

Also, the 6WINDGate software is available with optimized support for industry-leading multicore processors (currently, those based on Cavium, Freescale, Intel and NetLogic architectures, with others planned). The 6WINDGate architecture is modular, enabling the software to be efficiently ported (by 6WIND) to new architectures and with the hardware dependencies restricted to a limited number of sections in the overall code.



When 6WINDGate is ported to a new processor architecture, optimizations are performed to ensure the best possible utilization of on-chip resources such as offload engines, security accelerators and other processor-specific functions designed to maximize the performance of specific protocols or algorithms. At the same time, 6WINDGate makes full use of the services provided in the processor suppliers' Multicore Execution Environments.

6WINDGate enables system developers to select whichever processor is most appropriate for a specific end product, knowing that their application software, through the 6WINDGate stack, will be able to extract the best performance from that platform. At the same time, there's no need for the system developers themselves to become experts on the details of the processor architecture, since 6WINDGate provides an abstraction layer and itself implements the necessary support for performance-oriented features.

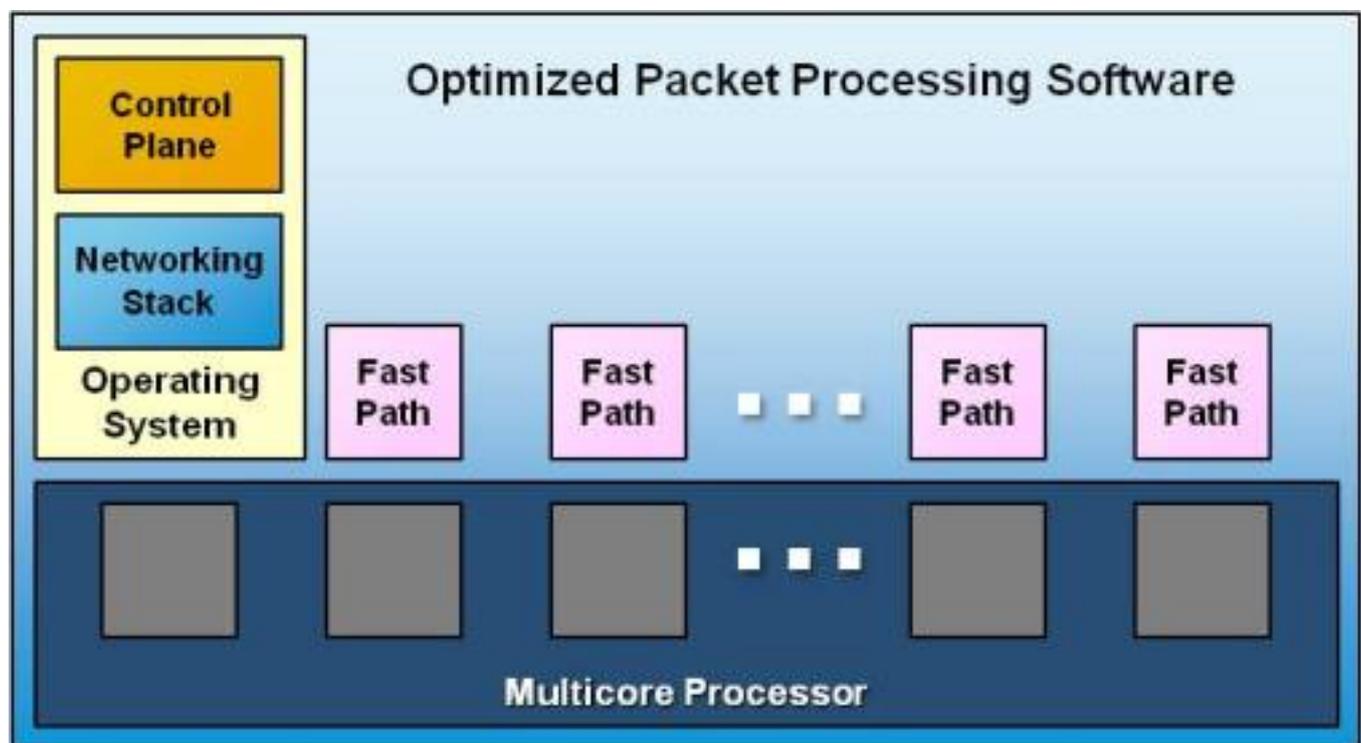
In terms of the performance challenges for next-generation networks, a standard networking stack uses services provided by the operating system and is subject to significant overheads associated with functions such as preemptions, threads, timers and locking. These processing overheads are imposed on each packet passing through the system, resulting in a major performance penalty for overall throughput. Furthermore, although some improvements can be made to an operating system stack to support multicore architectures, performance fails to

## Adopting Multicore Processors for 40G Telecom Equipment

Published on Electronic Component News (<http://www.ecnmag.com>)

scale linearly over multiple cores for complex packet processing such as required by 4G and a processor with, for example, eight cores may not process packets significantly faster than one with two cores for GTPu-to-GRE encapsulations. All in all, a standard operating system stack does a poor job of exploiting the potential packet processing performance of a multicore processor.

A superior solution is provided by specialized packet processing software such as 6WINDGate, optimized for multicore architectures. The networking stack is split into two layers. The lower layer, typically called the fast path, processes the majority of incoming packets outside the operating system environment and without incurring any of the operating system overheads that degrade overall performance. Only those rare packets that require complex processing are forwarded to the operating system networking stack, which performs the necessary management, signaling and control functions.



A multicore processor is well-suited to implementing this kind of software architecture. Most of the cores can be dedicated to running the fast path, in order to maximize the overall throughput of the system, while only one core is required to run the operating system, the operating system networking stack and the application's control plane.

Until recently, the only restriction when configuring the platform was that, since the cores running the fast path were running outside the operating system, they had to be dedicated exclusively to the fast path and not shared with other software. With the recent evolution towards a hybrid fast path model, the system can now be reconfigured dynamically as traffic patterns change in order to share the CPU

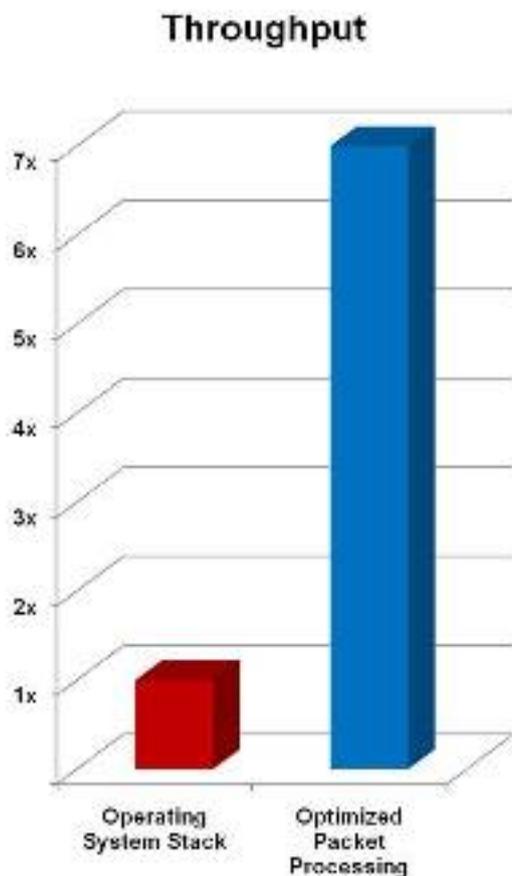
## Adopting Multicore Processors for 40G Telecom Equipment

Published on Electronic Component News (<http://www.ecnmag.com>)

---

resources allocated to the control plane and the fast path.

In a typical 4G application such as a packet gateway or switching gateway, when the standard operating system networking stacks are replaced by optimized packet processing software based on the fast path concept, the networking performance of the processor subsystems will typically increase by seven to ten times. This allows the TEM to meet system throughput goals that may have been unachievable on a single multicore processor when using a standard operating system stack.

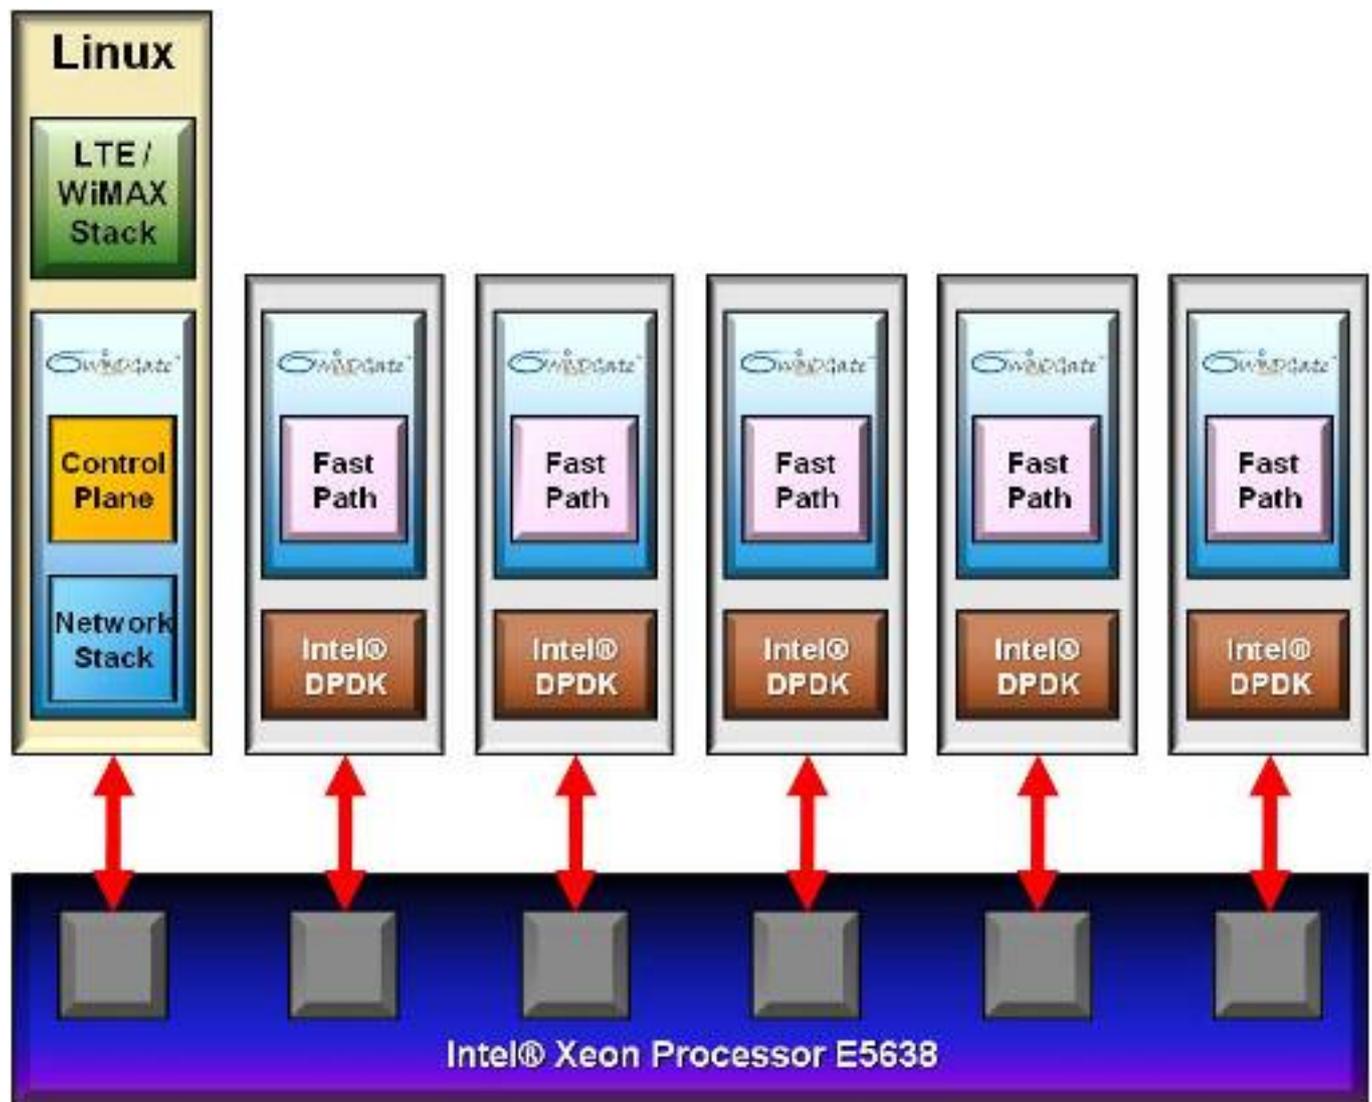


With a comprehensive set of protocols available for the control plane, the networking stack and the fast path, 6WINDGate provides developers with a single-vendor solution for all the protocols required for a high-performance wireless infrastructure platform based on multicore technology. By removing the need for developers to integrate networking software components from multiple suppliers, 6WINDGate has been proven to accelerate the time-to-market for networking equipment by up to twelve months.

When installed on an Intel® Architecture-based platform such as the GE A10200, for example, 6WINDGate can be configured at run-time to make the optimum use of the number cores available.

# Adopting Multicore Processors for 40G Telecom Equipment

Published on Electronic Component News (<http://www.ecnmag.com>)



In the example shown above, the six-core Intel® Xeon® processor E5638 is used as follows:

- One core is configured to run Linux and the application stack, as well as the control plane and the networking stack;
- The remaining five cores are configured to run the fast path, which makes full use of processor-specific services provided by the Intel® DPK.

GE offers several multicore processor boards to which 6WINDGate has been ported. Since 6WINDGate already supports a large number of protocols, it greatly reduces the effort required to create a benchmarking setup. It takes less than an hour to configure and run an IPsec or IP Forwarding performance test and get realistic measurements that reflect what the real application can be expected to deliver.

Consider the following example. GE recently ported 6WINDGate to its A10200 dual Xeon® processor E5638 ATCA blade. IP Forwarding was used to measure packet IO performance on the dual 10Gbit/s Ethernet Fabric interface. For the test, 64Byte packets were used to create the most difficult workload on the CPU. IP Forwarding performance data was measured using both 6WINDGate and the standard Red Hat

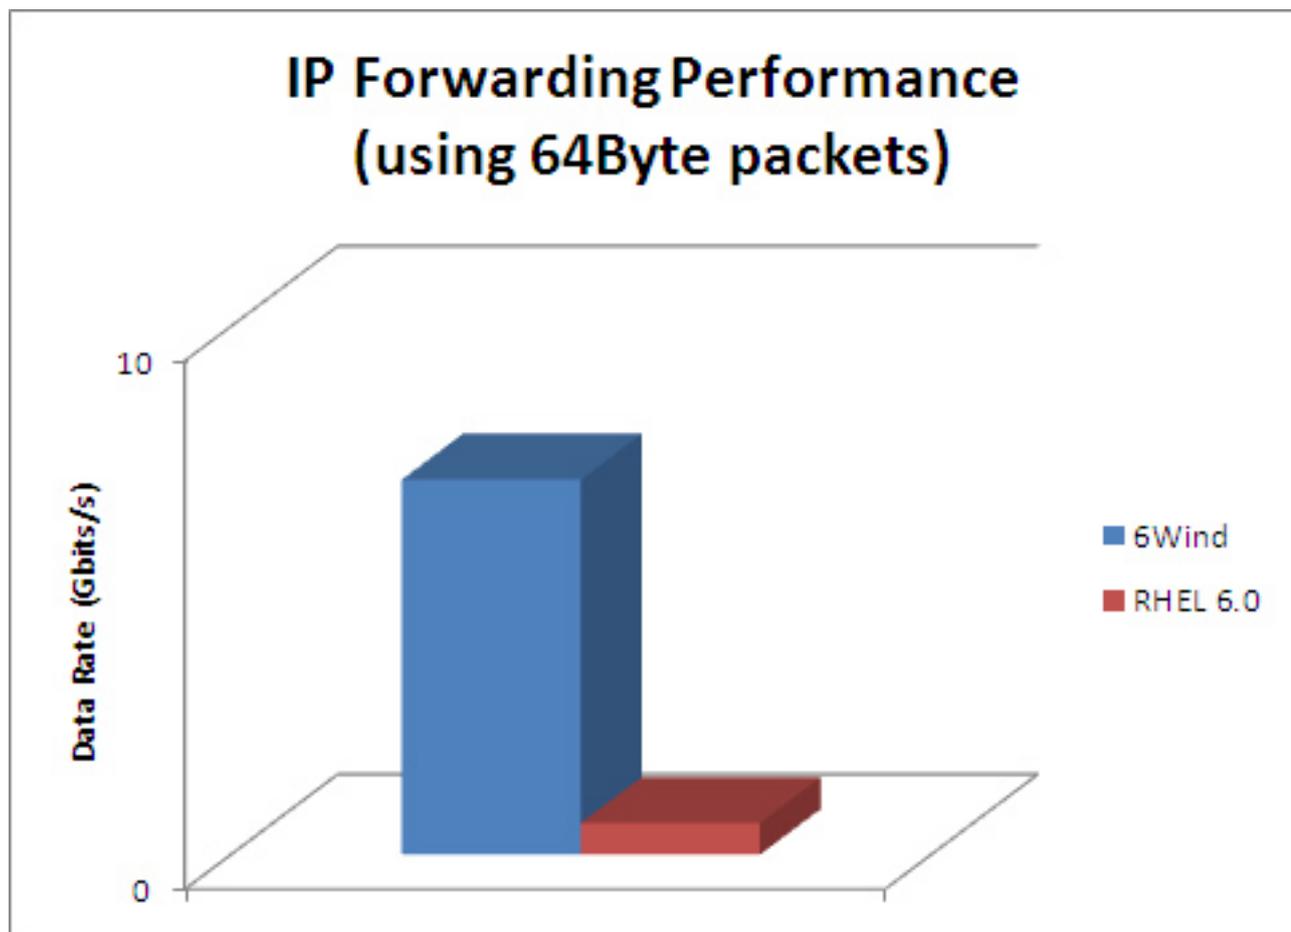
## Adopting Multicore Processors for 40G Telecom Equipment

Published on Electronic Component News (<http://www.ecnmag.com>)

---

Enterprise Linux 6.0 operating system. The resulting huge performance gap between the standard Linux operating system and the 6WINDGate software is due to the fact that 6WINDGate leverages Intel's latest bare metal DPDK software. This enabled programmers to dedicate some processor cores for specific packet processing tasks, such as IP Forwarding.

This example illustrates the challenge of how the same hardware can deliver vastly different performance depending on how well the available processor resources are utilized. Without proper performance testing, customers are left with speculative or even inaccurate estimates that might result either in performance that is less than expected, or in a significant over provisioning of compute resource and the associated extra cost.



6WINDGate has been ported to GE's Cavium OCTEON®-based and Intel® Xeon® processor E5638 ATCA blades, and GE expects to also have the software ported to upcoming products. Consequently, customers can run performance tests comparing different hardware architectures, while easily migrating from one hardware architecture to another and reducing dependency on the hardware specifics. This mitigates processor availability issues and future roadmap risks.

## Adopting Multicore Processors for 40G Telecom Equipment

Published on Electronic Component News (<http://www.ecnmag.com>)

---



Finally, when adding more than one supplier into the product equation technical support can be a challenge. GE and 6WIND have been working together for several years and engineers at both companies are very familiar with the other's products. This close working relationship reduces technical support challenges and accelerates the customer's application development process.

Driven by users and Service Providers, the Telecommunications industry is aggressively moving towards 40Gbit/s data rates in ATCA systems. However, packet processing at these high data rates is only feasible by using hardware-specific acceleration features and fully utilizing all available processor cores by using bare metal operating systems. Such architectures pose a number of significant risks in terms of performance estimation, silicon availability and application software migration. GE's ATCA blades coupled with the 6WINDGate software mitigate these risks by enabling customers to run accurate performance tests, reduce hardware-specific programming effort, minimize time-to-market, and enable painless migration of application software from one processor family to another.

**Source URL (retrieved on 07/07/2015 - 1:46pm):**

<http://www.ecnmag.com/articles/2011/05/adopting-multicore-processors-40g-telecom-equipment>