WHITE PAPER
Intel® Processors
Software Defined Networking

Software Defined Networking and Software-based Services with Intel® Processors

To remain competitive, today's network service providers need to be able to respond to evolving markets and traffic types.

Fixed line network service providers are coming under increasing commercial pressure from the growth of over-the-top (OTT) content providers. These new players are starting to identify individual end users, profile services for those users' specific needs, and establish direct billing relationships that threaten existing business models. Meanwhile mobile broadband service providers are facing a ten-fold increase in traffic over the period 2011-16.[1] This growth is coming from video-based services, which have a very low revenue potential per bit carried compared to the applications, such as Web browsing, that established the market pricing for the underlying network service.[2]

To remain competitive, today's network service providers need to be able to respond to evolving markets and traffic types in a timeframe of hours and days rather than the months and years typically involved in developing carrier grade network services. Recent improvements in Intel® microarchitecture open the door for network service providers to gain unprecedented flexibility and control over customer offerings through use of software-defined networks (SDN). By virtualizing network functions on Intel® architecture, network service providers can employ techniques such as deep packet inspection (DPI), geographic load balancing, and power management to optimize available bandwidth—resulting in dramatic cost savings. Moreover, network virtualization enables network service providers to quickly deploy software applications, enabling new revenue-generating services that can be brought to market in shorter timescales.

This paper describes how Intel® processor-based devices deployed at the network edge can help service provider network strategists, architects, and designers reduce initial development costs, on-going operational costs, and time to market for new services. It presents the benefits of using standard hardware for network components and describes how Intel® architecture-based hardware can fit into existing network environments.

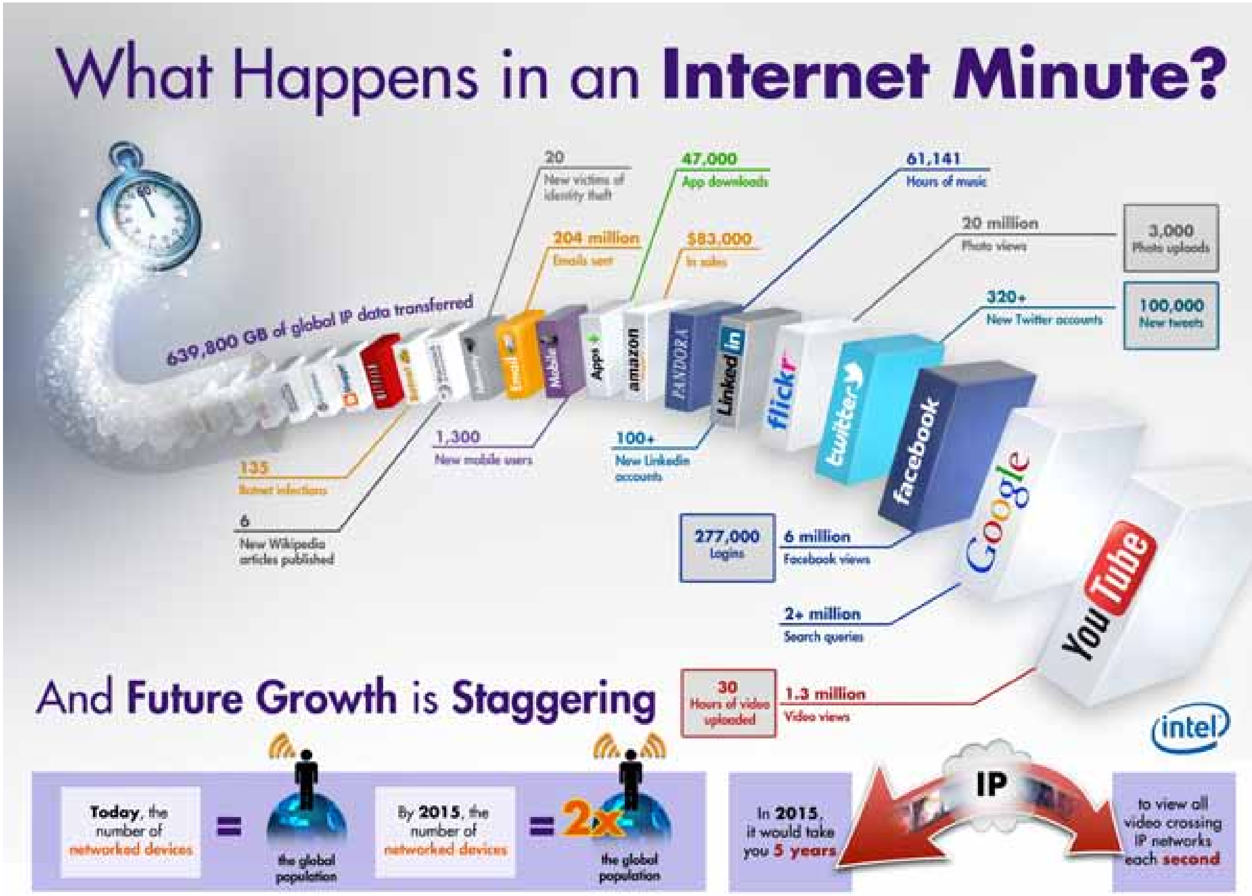

Intel® microarchitecture improvements have significantly reduced the need for specialist silicon, enabling network service providers to leverage the proven scalability of modern virtualized data center technology.

Network Virtualization: The Next Evolution of Service Provider Networks

Historically, network service provider business models were based on the ability to securely partition bandwidth on a shared platform between multiple circuits for individual customer services. Next-generation network architectures have enabled consolidation at the frame and packet level using IP, MPLS, and Ethernet technologies. However these approaches have had limited impact at the service level, which is the key to future revenue yield. To add services to current networks, network service providers generally add fixed-function packet processing appliances, creating a network that is difficult to power-manage, costly to scale, and cumbersome to expand.

To regain profitability, network service providers must either reduce network costs or find new means of generating revenue—or both. By deploying network functions as software applications, network service providers can gain flexibility lacking in current network environments and can achieve cost reduction and service velocity already proven in enterprise data centers.

Intel® Architecture: A Flexible, Scalable Approach to Network Expansion

Intel® architecture provides network service providers with a standard, reusable, shared platform for SDN that is easy to upgrade and maintain. Recent Intel® microarchitecture improvements have significantly reduced the need for specialist silicon, enabling network service providers to leverage the proven scalability of modern virtualized data center technology. Advantages of this approach include a streamlined network and cost savings through hardware reusability and power reductions. The adoption of Intel® processor-based network components also enables network service providers to integrate solutions currently being developed by independent software vendors (ISVs).

Microarchitecture Improvements Enable Carrier Grade Performance

Intel has made several improvements to the latest revision of its microarchitecture with with telecom and packet networks in mind. In addition to handling application and control plane processing, new Intel processors can perform packet processing and signal processing—workloads that have historically required specialist silicon devices and associated development tools to handle the level of traffic generated on communications networks. In addition, Intel platforms will offer two key security features for network service providers. Intel® Virtualization Technology (Intel® VT) provides secure, separate partitioning of processing capacity to prevent interference between network services. Similarly, Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI) functions can be used to accelerate encryption of control plane protocols and to create Virtual Private Network (VPN) capabilities on the underlying infrastructure—enabling the two functions to share the underlying processing capability.

In addition to Intel microarchitecture improvements, significant software optimizations are possible with Intel® Data Plane Development Kit (Intel® DPDK). Intel DPDK is a set of drivers and libraries provided to the industry to further improve high performance packet processing up to 10 million packets per second per core on Intel's latest generation processors. Furthermore, Intel VT also provides hardware assists to industry hypervisors and significantly reduces overhead compared to previous microarchitecture. While overheads of 30 percent were typical only two processor generations ago, the latest Intel VT can reduce this to less than 10 percent.

As a result, network service providers can now virtualize network workloads and consolidate application, control plane, packet processing, and digital signal processing workloads onto a single server. For example, service edge functions can be shared with control plane signaling requirements, such as SIP Call State Control Functions (x-CSCF) and Policy Charging and Rules Functions (PCRF). Similarly, network management and business support functions can be consolidated on the same shared infrastructure platform, or on dedicated instances of the same technology as required.

Standard Hardware Gives Network Service Providers More Control

By virtualizing network functions as software applications, network service providers gain flexibility in network configuration, enabling significant benefits, including cost savings and faster time to market for new services.

Flexible deployment of network functions. Intel architecture features a standardized instruction set that allows developers to create SDN elements that can be run across the portfolio of Intel processors, exploiting hardware functions as available. By using network components based on Intel architecture, network service providers can flexibly and consistently deploy network functions at various network locations—from the largest data center, to network equipment buildings using Intel® Xeon® or Intel® Core™ processors, down to street cabinets and customer premises equipment (such as home gateways and set-top boxes) using Intel Core or Intel® Atom™ processors. This flexibility allows network service providers to optimize deployment of network functionality to suit business needs or adapt available capacity for operational expediency.

Faster time to market for new services. The consolidation of packet and signal processing into Intel architecture means that servers based on the latest Intel processors can run these software-defined products without specialist silicon-based hardware. As a result, network service providers can evaluate and exploit software applications from the vast Intel ISV ecosystem, without acquiring and deploying specific hardware—significantly reducing time to market.

Hardware reusability. Once workloads are virtualized, network service providers can continuously optimize network usage by moving applications and services to where their deployment is most cost-efficient—without making fundamental architectural changes. For example, a new service might run on centralized compute infrastructure when first introduced and, once it has been proven in the market, be distributed to an edge-based infrastructure for high capacity revenue generation. The use of standard hardware also enables network service providers to repurpose installed infrastructure capacity to support different traffic demands. For instance, Intel architecture-based infrastructure can serve different purposes over the natural evolution of a service requirement throughout its lifecycle, or be repurposed on a time-of-day basis to meet peak hour traffic demands and undertake other processing requirements in off-peak periods.

INTEL® TOOLS ENABLE PERFORMANCE IMPROVEMENTS

Intel® Data Plane Development Kit (Intel® DPDK) – Special libraries and optimized drivers for enhanced I/O performance, allowing individual cores to communicate through Network Interface Controllers (NIC) at up to 10 Gbit/s

Intel® Signal Processing Development Kit (Intel® SPDK) – Tools to exploit the Advanced Vector Extensions (AVX) instructions and the on-chip graphics processors in 2nd and 3rd generation Intel® Core™ processors

Intel® Virtualization Technology (Intel® VT) – Hardware assists built into the latest Intel Core processors to increase the performance and security of virtualized applications

HANDLING VIDEO TRAFFIC DEMANDS WITH DEEP PACKET INSPECTION

Deep Packet Inspection (DPI) technology enables individual users' traffic to be classified and then processed according to various criteria, including fair sharing of available bandwidth. With DPI, network service providers can introduce intelligent content processing, such as local caching of popular content, transcoding source material into lower bit-rate coding standards that suit the terminal device, and transrating content to use the lowest number of screen pixels and frames per second necessary to still provide a good end user experience on the small screens of smart phones.

Similarly, network service providers can use DPI classification by traffic signature to create premium service packages that can guarantee end users a certain quality of experience during periods of network congestion.

Network Virtualization Helps Reduce Operating Costs

Network service providers are under relentless pressure to reduce operational costs, and are increasingly aware of the total cost of ownership (TCO) for their platforms. Decoupling services software from the underlying hardware introduces a new dimension for cost savings. In addition to reducing volumes of hardware and the associated energy costs, network service providers can exploit tools and technology proven in enterprise data center operations to decrease ongoing costs in line with traffic demands. Specific benefits include:

- **Reduced power consumption.** The ability to share a processor between multiple workloads enables servers to be loaded to levels based on their available capacity, rather than having to deploy specific devices that could be highly underutilized. Similarly, having a standard application environment means that instead of requiring dedicated production and backup appliances for each workload, infrastructure provided for standby resilience can be shared. This reduction in hardware leads to a corresponding reduction in accommodation and associated power consumption. In addition, Intel processors feature advanced power-management functions at the chip level that can power down unused cores within a processor—delivering further power savings.
- **Simplified field maintenance.** The standardization of a single server model reduces spares holdings and introduces the possibility of consolidation with cloud computing or other data center support arrangements. Furthermore, application software maintenance releases can generally be downloaded to a remote server—unlike firmware upgrades, which often require a field upgrade with associated truck rolls and many combinations of hardware, firmware, and software to handle, support and track.
- **Lower-risk new service trials.** New service trials can be instantiated on existing servers rather than having to deploy new devices for an unproven service requirement. If market take-up does not meet expectations, the underlying hardware can be reused to provide capacity for other services.
- **Improved responsiveness to service demands.** With standard hardware, network service providers can move some workloads around the network to tactically handle capacity issues pending infrastructure rollout in specific geographical areas.

One Platform, Many Applications

The latest Intel processors are suitable for a wide range of applications in both fixed and mobile next generation networks, including:

- At the service edge, where streams of IP packets can be recognized and processed as voice, video, or other data services
- As access gateways—such as the Broadband Network Gateway (BNG), the Generalized GPRS Services Node (GGSN), and the Packet Data Gateway (PGW)—with tunnel termination, QoS handling, and authentication interfaces
- In security devices that enable safe attachment of service provider networks to the Internet, including firewalls, content filters, and denial of service attack filters
- Within the mobile core, as well as in radio access network (RAN) devices, such as radio network controllers (RNCs). Work has also begun on development of commercial 3GPP long term evolution (LTE) base stations based on Intel® processors.[3]

Putting Intel® Architecture to Work in the Network

The transition from fixed-function network appliances based on specialist silicon to Intel processor-based devices is straightforward. Intel processor-based network infrastructure elements can be designed to be physically interchangeable with their fixed-function counterparts and to have familiar network service provider interfaces. The difference is on the inside. Each Intel architecture-based network device has three key components that cooperate to form a complete system:

- **Processor.** The latest Intel processors feature up to 10 independent processing cores running at clock speeds of up to 3 GHz and include technologies—such as shared and secure cache memory, advanced power management, and Intel® Hyper-Threading Technology, which effectively doubles the number of virtual cores—to achieve the best performance for dynamic workloads.
- **Interfaces**. The processor is connected to the outside world by a variety of components, typically presenting PCIe* and USB interfaces. Ethernet interfaces, including Intel® 82599 10 Gigabit Ethernet controllers, are designed specifically to handle virtualized networking over multiple 10 Gbit/s links. Additionally, Intel offers high-performance Solid-State Drives (SSDs) with capacities up to 600 GB packaged with industry standard interfaces, enabling complete systems with no moving parts for improved overall reliability.
- **Form factor.** Intel devices can be organized to fit a variety of form factors depending on the power and performance required. A common design uses two processor sockets, enabling installation of appropriate multi-core devices, along with memory, dual 10 Gbit/s Ethernet interfaces, and a number of standard disc drives. This device configuration is generally available as a boxed rack-mount server mounting in one unit of a 19-inch rack. Alternative packaging includes blades that can be slotted into chassis enclosures, typically meeting the Advanced Telecom Computing Architecture* (ATCA*) format or an equivalent from a telecom equipment manufacturer or an IT vendor, with alignment to NEBS environmental standards.

For specific high performance computing requirements, up to eight Intel architecture-based multi-core devices can be interconnected within one processor complex, as in the HP* Integrity* BL890c i2* server blade.

Intel also works with operating system vendors to ensure that their products can take full advantage of Intel architecture to support operational workloads. As a result, a number of operating systems are optimized for Intel architecture.

The choice of operating system depends on higher level software application needs. With Intel architecture-based infrastructure, network service providers can quickly repurpose hardware by loading a different operating system. This gives network service providers the flexibility to optimize utilization of available infrastructure on a time-of-day basis, or to potentially repurpose development systems for production use for business continuity in the event of a major incident.

Intel® Architecture Provides Choice of Operating Environments

- **Hypervisor.** A hypervisor enables applications to operate in a virtual machine environment running their own local operating systems, giving network service providers maximum flexibility. Examples include VMware* ESX*, Microsoft* Hyper-V*, and RedHat* KVM*.
- **Real-time operating system.** Real-time operating systems such as Wind River* VxWorks* deliver optimized communications performance.
- **Asymmetric operating system.** These hybrid operating systems allow some cores to act as real-time machines for high performance communications applications, while other cores can support multiple environments for control plane and management functions. Examples include Wind River*Network Acceleration Platform* and 6Wind 6Windgate*.

A Trusted Partner, Today and Tomorrow

Intel has been anticipating the shift to SDN and is committed to helping network service providers take advantage of this opportunity. Intel developer tools, such as Intel DPDK and Intel® Signal Processing Development Kit (Intel® SPDK), enable the global Intel ISV community to easily develop optimized software and services for telecom networks. Furthermore, the roadmap for the next two generations of Intel microarchitecture indicates that these versions will continue to progress according to Moore's Law, increasing the number of cores and performance available in a single device. This is underpinned by Intel's new 3D transistor technology,[4] which has opened the door to 14 nm and 10 nm devices that will deliver even more product performance in the future. By adopting Intel architecture-based network components, network service providers gain access to a continually expanding body of optimized software applications and increasingly higher performance hardware components for years to come.

Planning Your Evolution to Intel® Architecture

Evolution of service provider networks to SDN has been predicted for several years. The time has come for network service providers to plan their transition for competitive advantage. Thanks to recent microarchitecture developments, Intel processors now deliver the packet processing performance needed to support the growing demand for service edge capabilities. By virtualizing these workloads on Intel architecture-based edge infrastructure, network service providers can realize significant reductions in operating expenses while gaining the flexibility and scalability to compete in today's rapidly changing telecommunications market.

In addition to working with equipment manufacturers, Intel is already exploring the benefits of SDN with several leading fixed and mobile network service providers in Europe, the Americas, and the Far East. In October 2011, Verizon revealed their intentions in public,[5] and a number of network service providers and equipment vendors are expected to make similar announcements during 2012. Furthermore, Intel is working with Hewlett Packard* and Verizon* to test and validate a number of network functions on HP* ProLiant* servers using Intel processors, and to determine how cloud concepts and SDN techniques can best be applied to fixed and wireless networks. Look for more information on these exciting projects in the near future.

The first step toward planning your evolution to network virtualization is to consider the impact of software-based approaches on your network. As you consider changes to your network design, ask your suppliers how they expect to handle the transition to open, cost-effective infrastructure that will support your service innovation needs. In the meantime, consider specifying Intel architecture-based equipment for your current network and service appliance requirements to facilitate future evolution to an SDN.

For more information about Intel architecture, visit http://www.intel.com/content/www/us/en/communications/intel-service-edge-animation.html

[1] Various speakers, BroadBand Traffic Management Congress, London, November '11

[2] Mark Roettgering, T-Mobile USA, eComm Conference, '09

[3] Justin Rattner, Intel CTO keynote, Intel Developer Forum September 2011

[4] "Ivy Bridge" announcement, April 2011

[5] Stu Elby, VP Network Architecture, Verizon, Open Networking Summit, October 2011

Printed in USA 0212/MS/SD/PDF Please Recycle 326846-001US