WHITE PAPER

Communications
Data Center Solutions

Building the Networks of the Future: The Rise of the Scalable Data Plane

The networks of tomorrow will rely on an increasingly intelligent, programmable and reactive data plane, running on industry standard network hardware, an open source software ecosystem, and Intel® architecture, including processors, field programmable gate arrays (FPGAs), network interface cards (NICs), and more.

Data Plane Requirements

1. Intelligent
2. Programmable
3. Reactive
4. Secure
5. Performant

Executive Summary

- As applications move increasingly to the cloud, cloud and network infrastructure technologies are converging. Network environments are becoming software-defined and virtualized—creating both disruption and opportunity in the marketplace.
- Communications Service Providers (CoSPs) are being forced to rearchitect their networks to be more scalable, programmable, and flexible at the edge to rapidly innovate, deploy new services, and protect customer data from malicious attacks.
- Next-generation networks must be robust, service-aware, and cost-effective. Additionally, they must be able to support the needs of the coming "cloud-native" applications that will operate independently and unaware of infrastructure.
- Born in the cloud, these applications will rely on the availability of secure, scalable network architecture and an intelligent, reactive and performant data plane. Traditional data plane infrastructure, including purpose-built ASICs and FPGAs, can't meet all of these requirements.
- As the separation between hardware and software continues, data planes will need to evolve to incorporate industry-standard network hardware, an open source software ecosystem accessible through open APIs, and Intel® architecture.

Table of Contents

Introduction

Convergence, separation – the network environment is changing dramatically, elevating the importance of a scalable, flexible, and performant data plane in processing and delivering application data. Cloud and network computing, enterprise and cloud infrastructure, fixed and mobile are all converging and, at the same time, hardware and software are decoupling. Tectonic shifts are happening in communications, as the industry prepares itself for a wave of non-traditional data traffic generated by the Internet of Things, 5G use cases and services, and a new generation of "cloud-native" applications born and residing exclusively within the cloud.

These applications will have a heightened dependence on packet data, requiring Communications Service Providers (CoSPs) to consider and plan for an enhanced data plane with next-generation intelligence and processing capabilities. Traditionally, Intel architecture has been used to manage the control plane on physical appliances, while purpose-built ASICs have been used to manage the data plane. ASICs have a defined role to play in areas of the network requiring fixed functions but offer less flexibility at a higher cost when it comes to delivering on CoSPs' need for agile service provisioning at the edge. Because of its flexibility, Intel architecture has an increasingly vital role in addressing both control and data plane processing needs. Data plane acceleration technologies, such as Intel® QuickAssist Technology (Intel® QAT) and Field Programmable Gate Arrays (FPGAs) also offer tangible benefits in these areas.

At Intel, we believe four tenets are foundational to the work of building an efficient, programmable and scalable data plane: subscriber-centric computing, industry-standard server infrastructure, software architecture consistency, and dynamic network security. Achieving a next-generation data plane will require more than a revision of the CPU complex—we will need to reimagine all aspects of how the processor and platform are architected and administrated, combining industry-standard servers, Intel architecture, and open source software for optimal results. These changes will enable CoSPs to improve the speed of innovation, product development, and service delivery.

THE 4 TENETS OF AN EFFICIENT, SCALABLE DATA PLANE

- Subscriber-Centric Computing
- Industry-Standard Servers
- Consistent Software Architecture
- Dynamic Network Security

The purpose of this paper is to outline the benefits of data plane processing on Intel architecture, and explore how Intel is addressing the data plane today while also driving towards our vision of an efficient, programmable and scalable network of tomorrow.

Networks Built Around Subscribers

CoSPs began the process of rearchitecting their networks back in 2012, leading the foundation of the European Telecommunications Standards Institute (ETSI) as software-defined networking (SDN) and network functions virtualization (NFV) began to gain steam. Today, they are evaluating options for re-architecting the edge of their networks with data center technologies to meet increased demand from subscribers and a growing number of machine-to-machine applications.

Mobile data traffic is growing at a CAGR of 46 percent[1], forcing CoSPs to reevaluate how they deliver both enterprise and consumer services—finding new approaches to controlling costs and build-out requirements, while also developing (and growing) additional services and revenue streams. CoSPs are increasingly challenged to find ways to meet the emerging and often divergent needs of both residential and enterprise customer segments with the same network infrastructure.

Today, the provisioning of consumer data, voice, and video and enterprise VPN services involves the configuration of different types of service functions on an edge router. The Broadband Network Gateway (BNG) provides subscriber broadband internet services while the Provider Edge Router typically connects the enterprises, providing them with global VPN connectivity through the Multiprotocol Label Switching (MPLS) network. The move to edge-hosted VNFs gives CoSPs a more flexible and cost-effective approach to provisioning such services. By adding aggregation points close to the user, CoSPs can avoid backhaul traffic and reduce latency restrictions, which both simplifies traffic and reduces costs.

Mobile data traffic is growing at a CAGR of 46 percent, forcing CoSPs to reevaluate how they deliver services.[1]

This also allows CoSPs to customize services for disparate markets. For example, in rural non-commercial regions, it makes sense to deploy a Virtualized BNG (vBNG) function for internet subscriber management, while in urban or commercial centers the vBNG and a Virtual Provider Edge (vPE) function are required to provide the requisite enterprise services. This service flexibility has significant advantages. Still, at the core data center, the traditional bare metal servers, and even the virtual machines, are creating a software maintenance load that is not scalable.

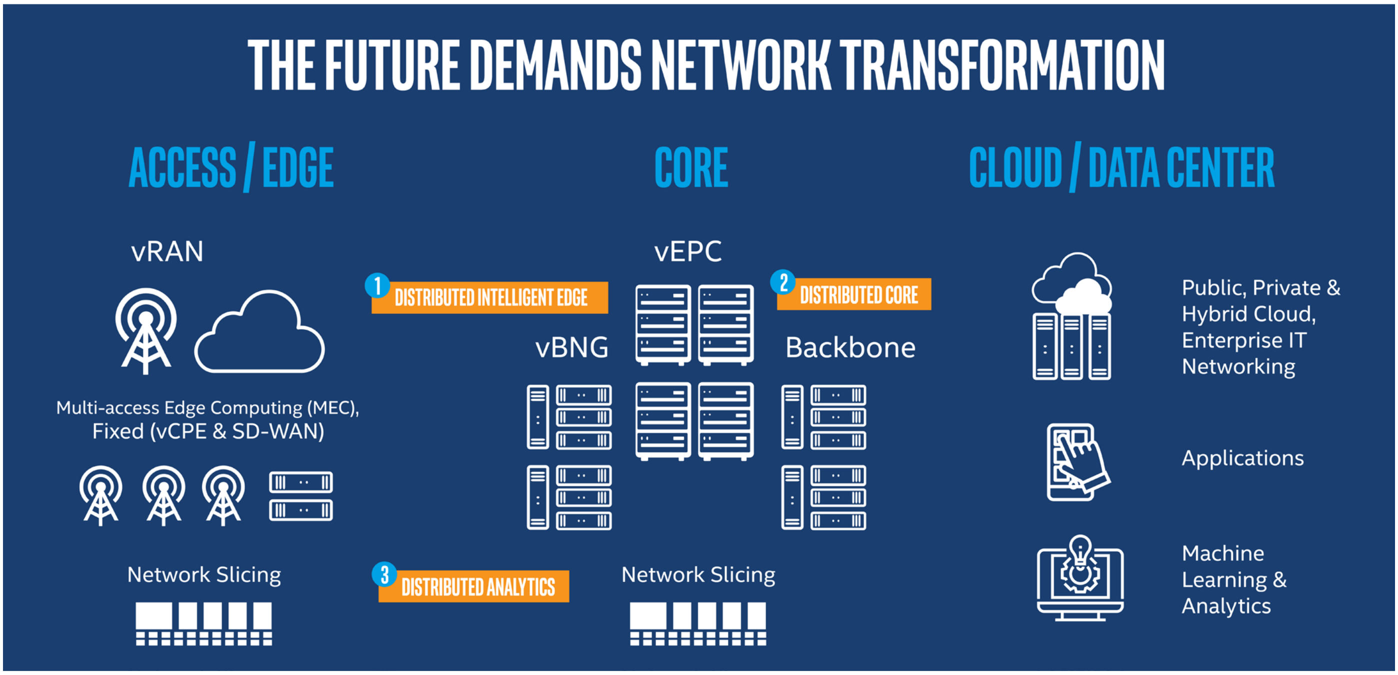
THE FUTURE DEMANDS NETWORK TRANSFORMATION
ACCESS / EDGE
CORE
CLOUD / DATA CENTER
vRAN
Multi-access Edge Computing (MEC), Fixed (vCPE & SD-WAN)
Network Slicing
1 DISTRIBUTED INTELLIGENT EDGE
vBNG
3 DISTRIBUTED ANALYTICS
vEPC
2 DISTRIBUTED CORE
Backbone
Network Slicing
Public, Private & Hybrid Cloud, Enterprise IT Networking
Applications
Machine Learning & Analytics

Bandwidth demands on compute nodes are exceeding traditional vSwitch capabilities, and alternatives are being pursued to free up computing cores to manage the load. CoSPs need new operating models for cloud-ready infrastructure.

For CoSPs, many of the NFV use cases identified refer to functions and services which can be hosted in their IP Edge. For example, VNF-as-a-Service (Virtual Enterprise CPE), Virtualization of CDNs (vCDN), Virtualization of the Home Environment (Virtual Consumer CPE) and Virtualization of the Fixed Access Network (vBNG). The SDN use cases in the data center primarily focus on the rapid and flexible reconfiguration of the network planes to accommodate a software-defined infrastructure and storage with encapsulation using vLAN and vxLAN tunnels and vSwitch implementations on virtualized general-purpose compute nodes.

These use cases are driving innovation in instances where there is a locational advantage to the CoSP hosting the relevant VNF service in the IP edge close to the end user, and in the core data center on the compute nodes. Infrastructure at IP edge, central office environments, and core data centers is being transformed to enable that advantage with flexible service delivery.

Infrastructure Built to Industry Standards

> **Next-generation services require a next-generation data plane. Fixed function hardware doesn't lend itself well to today's level of diversity and rate of change.**

Next-generation services require a next-generation data plane, built with architecture designed to deliver on a wide variety of networking needs at ever-increasing data rates. Traditional fixed-function hardware infrastructure doesn't lend itself well to this level of diversity and rate of change. Industry-standard servers with software-based microprocessors are better suited for rapid deployments and a subscriber-centric network. With an industry-standard server—including network I/O, acceleration and the architectural consistency inherent in that—CoSPs can better maximize their investment in infrastructure to meet the needs of a highly-dynamic marketplace.

Today's multi-service edge routing appliances are typically engineered using chassis housing ASIC or merchant NPU line card implementations. These line cards are connected and switched using proprietary backplane technologies, providing exceptionally high terabit-capable solutions. Data centers are traditionally built out using multiple big iron switches that create a folded Clos network, which limits their ability to scale with user and service growth. In the current scenario, software upgrades and service innovation, which are directly tied to the hardware, typically occur in long release cycles of 18 to 24 months. This time-restricted approach places service providers at a significant commercial disadvantage.

Existing options for service innovation just aren't sufficient. Proprietary ASICs are somewhat expensive to design, test, and manufacture, and the development skill set (register-transfer layer design, silicon validation) required is in short supply. Considering that these development and manufacturing costs must still be recouped across a shortened product lifecycle, delivering new services can be slow and expensive. With that said, ASICs do have their place in next-generation networks, for fixed functions in the transport domain where the protocol stack (OSI layers 1-4) is well-understood, standardized, and aggregate switching speed is the premium technology choice factor.

Network Interface Cards (NICs)—including virtual NICs—play an essential role in moving network packets efficiently in and out of the platform. Typically, NICs provide stateless offload capabilities like packet filtering, TSO, Checksum, RSS and multi-queues for efficient distribution of traffic in a multi-core platform. With network virtualization, NICs added capabilities to handle overlay encapsulations like vXLAN and NVGRE, and to provide stateless offload functions to accelerate software-based virtual switching implementations in the host. This solution offers the right balance of flexible implementations in software with the right accelerations in hardware to achieve optimum performance.

With the rapid pace of innovation in the SDN and NFV market, the data plane must handle additional types of packets and encapsulations. Therefore, NICs are being further enhanced to provide programmable pipelines that can handle new frame formats for network applications, with support from technology such as Intel Dynamic Device Personalization (DPP). Network I/O performance is growing at a rapid pace to 100Gbps and beyond, which means software implementations need a boost in performance with hardware acceleration. This includes accelerating the virtual switching and efficient virtual interfaces for moving data to and from VMs or VNFs.

The industry has also seen the introduction of Smart NICs as an option to address additional packet processing or virtual switching needs. Because there is no industry standard for Smart NICs, different vendors are providing a range of solutions for vSwitch offload and other capabilities like Crypto/IPsec acceleration for secure transport of network virtualization overlay traffic. These first-generation solutions range from network processor-based implementations to programmable state machines to FPGA-based implementations. CoSPs must weigh the benefit of compute cycles savings from adding specialized hardware (e.g. Intel QAT accelerator, SmartNIC) to industry-standard high-volume servers against the additional operational and application certification complexities. CoSPs are waiting for acceleration from multiple vendors at a common application interface while also making available a software fallback so that applications can be developed to work across both accelerated and non-accelerated deployments.

Software for Every Application

As the industry adapts to the migration of applications to the cloud, it's also setting its sights on the coming cloud-native paradigm. Born from the advances in virtualization and cloud technology over the last decade, the term "cloud-native" refers to applications developed to only exist within and have all their dependencies satisfied by the cloud. It's

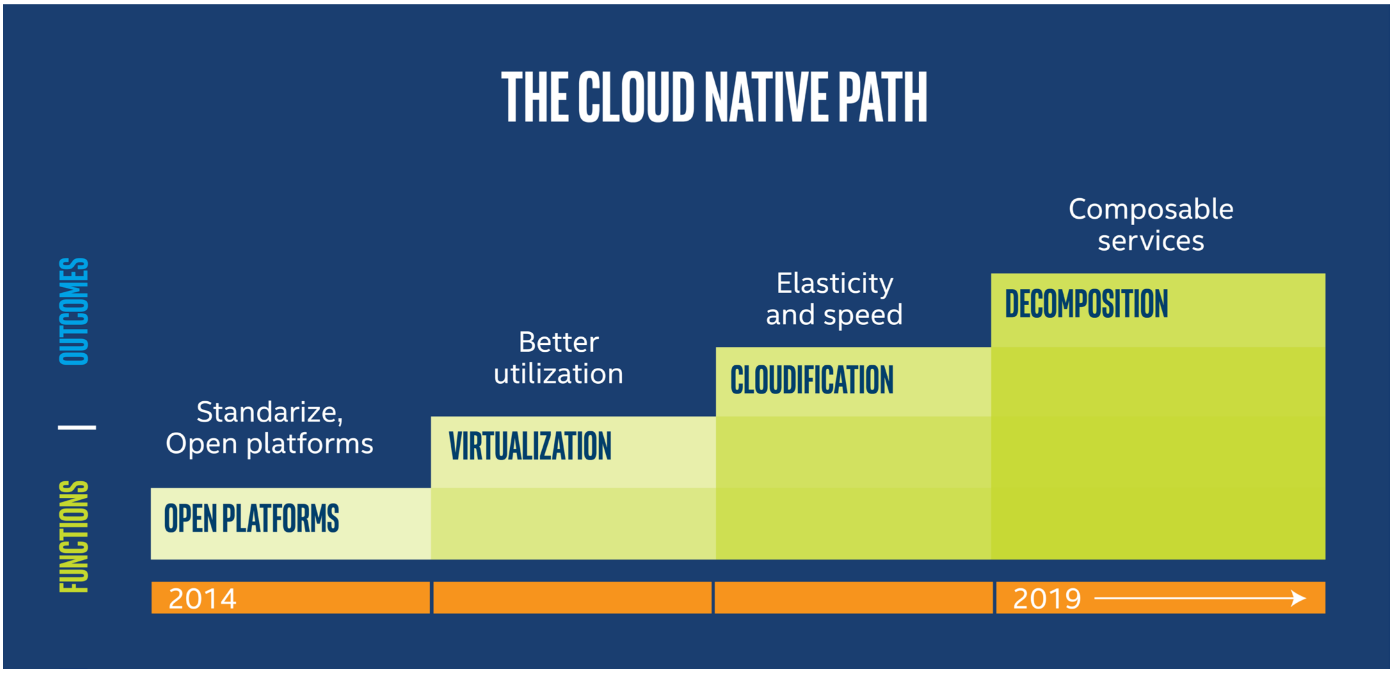

a shift from the physical or virtual machine in the cloud to the application in the cloud. This change will be reflected across all network domains, impacting the way services are architected, developed, deployed and managed. For example, environment management will be done per application—or even per application component—requiring scalable software designed to be distributed and reused across many types of new infrastructure, in a way that separates the application from the hardware. A resilient, stateless, self-healing micro-service architecture enables CoSPs to advance the industry toward cloud-native transformation.

Many CoSPs are currently considering cloud-native design, although it is not yet in the deployment cycle. Implementing a cloud-native system holds the possibility of an improved TCO, with rapid innovation fueled by a DevOps approach, the ability to scale up or down resource utilization, and OpEx/CapEx improvements. The cloudification of the network is in its very early transformation stage and requires re-architecting applications for the cloud with no dependency on network topology and the underlining hardware. Implementation can be achieved by packaging the software in VMs or containers, although containers are often the preferred method. With that said, there are organizational, operational and technical barriers that must be overcome for cloud-native to become a reality for CoSPs, particularly as it pertains to data plane processing.

Today's data planes are much more complicated than they were a few years ago, and connectivity services form only a part of the complete picture. Simple Layer 2/Layer 3 network services are necessary, but no longer sufficient to serve business or consumer needs. Currently, most Customer Premise Equipment (CPE) contains one or more connectivity options for Layer 1-3 including cable, fiber, Wi-Fi, and wireless, terminated in a Broadband Network Gateway. But CoSPs aren't just delivering connectivity anymore; they're also providing Layer 4-7 services, like IP-TV, to compete with OTT vendors such as Netflix*.

This means that tomorrow's cloud-native infrastructure will need to be enhanced to support per-subscriber SLAs via new network features and protocols. In the traditional model, this may require replacement of fixed function ASICs. In a programmable data plane, this means new software solutions. An open standard data plane, running on industry-standard hardware, can support new services without requiring new infrastructure as it is software-based.

An open-standard data plane, running on industry-standard hardware can support new services without requiring new infrastructure.

Network Security You Can Trust

The new role of the data plane isn't limited to delivering connectivity or IP-based services. It also includes ensuring that those connections and services are protected. Most medium- and high-end home routers today provide services like malicious website blocking, DNS filtering, and content filtering for mature content, infected device prevention and blocking. While they may seem optional today, these services won't be optional in future networks.

Botnet attacks are increasing as IoT devices expand. For example, the Mirai-botnet which targets Linux*-based connected devices (in this case home routers, IP cameras, DVRs, etc.) which appeared in September 2016 used IoT devices to mount a DDoS attack against Dyn. That attack disrupted services for Github*, Twitter*, and Netflix among others. In November 2016, Mirai was subsequently also the cause of an attack against 900,000 Deutsche Telekom* customers, interrupting service after infecting their DSL routers. DDoS attacks such as these cause not only colossal business outages but also consume mass quantities of network bandwidth.

Stopping DDoS attacks such as these, as well as more sophisticated "controllable" botnets that can be targeted on demand, requires every packet passing through a network

Network security relies on an intelligent data plane that encompasses much more than the L2-L3 layers of today.

to be examined against an ever-changing and evolving series of threats. Increasingly those packets are being passed through the data plane in the L4-L7 space, which means exploring them is well beyond the capability of ASICs—or for that matter any fixed programmable device. Fixed functions embedded in hardware do not provide a reactive platform. Network security relies on an intelligent data plane that encompasses much more than the L2-L3 data planes of today.

Along with preventing attacks, CoSPs are also concerned with providing robust service assurance to their customers. This is critical in the transformation to a software-defined and virtualized network environment. Service assurance in a virtualized world requires continuous monitoring of the platform hardware, the data plane, the environment, and software. The platform components and tools needed for the provisioning of resources include the collection of a growing set of platform performance, fault and other useful data, and the sharing of that data with management, analytics, and orchestration systems. This enables physical, virtual and service resources to be more thoroughly provisioned, managed, monitored and measured.

Intel Architecture: Supporting the Future of Business

Enabling next-generation data planes doesn't necessitate a choice between Intel architecture and custom silicon. Both are required. The challenge is to build future networks in a manner whereby functions constructed on Intel architecture and other technologies such as FPGAs and NICs complement each other to deliver service-aware networks that are performant, agile and cost-effective.

Intel architecture provides a range of platforms to address the need for a scalable, flexible, and intelligent data plane powered by open software on industry-standard servers. This includes solutions for lower-end router/packet processing applications with Intel Atom® processors; mid-range routers and appliances with the Intel® Xeon® D processors, and higher-end routers, security and network services platforms with Intel Xeon processors. With cloud and network virtualization, network functions are virtualized on standard servers, where a variety of virtualized network applications such as vRouters, vFW, vVPN, vPE, vCPE, vNAT, IDS/IPS, vLB can be deployed on demand.

Intel architecture is well-equipped to handle the challenges and requirements of the next-generation data plane. Intel platforms provide a rich set of virtualization capabilities, with the network I/O, encryption and hardware-assist and compression acceleration capabilities for both NFVi and applications to help address the network transformation challenges with SDN and NFV. This allows VNFs to be deployed unmodified at low, mid and large platforms permitting for reduced development costs and improved operator agility.

Intel's broad set of silicon product offerings are complemented with the Data Plane Development Kit (DPDK) Framework, which provides a standard methodology and a suite of libraries for building packet pipelines and complex packet processing applications for both traditional network appliance and virtualized network function usage models. The DPDK Framework can be used to build out NFVi infrastructure in the base platform or for developing virtualized network applications in support of cloud-native.

The effort led by the DPDK community to provide an application with a base set of network I/O access (e.g. native drivers, SR-IOV VF drivers, VirtIO drivers, VMXNET3) along with the orchestrators being able to deploy is quite mature. With that said, the DPDK community and broader orchestration communities are early on in establishing common approaches to both NFVi acceleration (e.g. OVS acceleration from a SmartNIC) as well as application acceleration (e.g. HQoS, IPSec, Public key encryption).

Multiple execution environments and application development for special purpose acceleration environments will always be a challenge. One of the key reasons for SDN and NFV network transformation is for customers to move away from specialized hardware environment to industry standard servers powered by flexible microprocessors like Intel architecture. With this combination of industry servers and architecture, software optimizations will continue to evolve to meet the business needs of CoSPs.

The challenge is to build future networks in a manners where functions constructed on both Intel architecture and other technologies complement one another—delivering service-aware networks that are performant, agile, and cost-effective.

[1] https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/complete-white-paper-c11-481360.html

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.