ETHERNITY NETWORKS

Enabling the Virtualized Edge with FPGA SmartNIC Data Acceleration

White Paper

ENET FPGA

Enabling the Virtualized Edge with FPGA SmartNIC Data Acceleration

Contents

Introduction: The Case for a Virtualized Edge

While the data explosion is now ubiquitous, there is an aspect to it that has yet to be fully addressed. The advent of 4K video, the Internet of Things (IoT), and augmented reality (AR), especially when accessible from mobile devices, has placed new and challenging demands on network operators to provide uninterrupted, high-quality service to an exponentially growing volume of devices and sensors for a massive number of end users.

For example, standard video has historically required 3 Mbps of throughput to enable streaming, which advanced to 10 Mbps high-definition video. Netflix recommends that an end user device be capable of receiving 25 Mbps to stream 4K ultra-high definition video[1]. Already in 2017, video accounted for 56% of mobile data traffic, with about 15 exabytes of data per month, and that is expected to grow to about 73% (107 exabytes per month) in 2023.[2]

To support the explosive growth in data usage, Communication Service Providers (CSPs) must upgrade their infrastructure with higher transmission capacity, such as Gigabit-to-the-Home and 5G mobile, while simultaneously bringing the data center infrastructure closer to user devices. Furthermore, CSPs are transitioning to replace rigid, vendor-locked legacy communications equipment at the edge of the network with Commercial-Off-the-Shelf (COTS) servers and bare metal switches to enable communications software appliances to run from multiple vendors on top of non-proprietary hardware, resulting in complete open networking. In so doing, they are reclaiming their freedom to select vendors and control the CSP roadmap for delivering better services to their customers.

1 Netflix Internet Connection Speed Recommendations: https://help.netflix.com/en/node/306

2 Ericsson Mobility Report, June 2018: https://www.ericsson.com/en/mobility-report/reports/june-2018

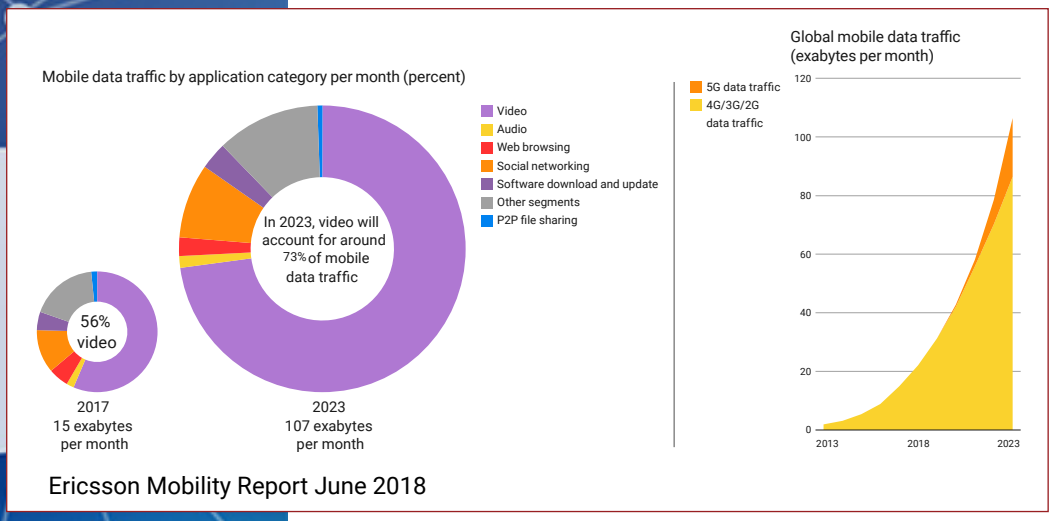

Figure 1:
The Rise of Video Across Mobile Traffic

The primary pressure point within the network, therefore, has shifted from the data center to the edge, which is defined as the area that distributes traffic originating from the carrier data centers and aggregates traffic from the end users, whether mobile, residential, or enterprise. To account for the greater demand for high performance networking all the way to the subscriber, the next generation network is relying heavily on 5G protocols, which are designed to bring high bandwidth and low latency to the network edge.

The primary pressure point within the network has shifted from the hyperscale data center to the edge.

In traditional communications networks, the network edge has usually taken the form of a point-of-presence location, which primarily distributed network traffic from a central office (CO) to a set number of subscribers. As data usage is growing per subscriber, though, the infrastructure must be upgraded to handle the larger capacity, and larger pipes are required to deliver so much content. To access servers at the point-of-presence, advanced networking features, such as routing, deep packet inspection (DPI), cybersecurity, and subscriber management, must then also be added to the edge. Furthermore, the content itself must be brought closer to the users. If these services and the content on which they are being applied are maintained only in the central data center, or even in the CO, the added latency becomes a major bottleneck.

This has led to the introduction of a virtualized edge of the network. To implement advanced features on open network hardware at the points-of-presence, CSPs are introducing Network Function Virtualization (NFV) into the edge, improving performance toward the end users and significantly reducing their operating expenses, such as power and physical footprint. They can also future-proof their capital investments as they avoid vendor lock-in and reduce the amount of equipment that requires regular updates.

Similarly, tier 1 central cloud providers, such as Amazon, Microsoft, and Google, want to improve their services and collect information more immediately by extending their cloud infrastructure closer to their users, bringing virtualization to the network edge. For example, an application such as Content Delivery Network (CDN) video

benefits greatly by having its cloud-based services located on the network edge instead of in a central data center, since it can more easily cache data and customize its cached contents to localized information. This move has also led to the evolution of Edge Computing, in which data is shared and computed across distributed devices instead of in the central cloud.

There are nearly a dozen separate initiatives to encourage standardization and enable open source virtualization at the edge, including CORD (Central Office Rearchitected as a Data Center); OPNFV's Edge Cloud, vCO, and Edge NFV projects; StarlingX; Akraino Edge Stack; EdgeX Foundry; ETSI's Multi-access Edge Computing (MEC) ecosystem; and the Open Edge Computing initiative. There are also efforts from all three major cloud providers to extend to the edge, in Azure IoT Edge, Lambda@Edge, and Amazon Cloud Front.

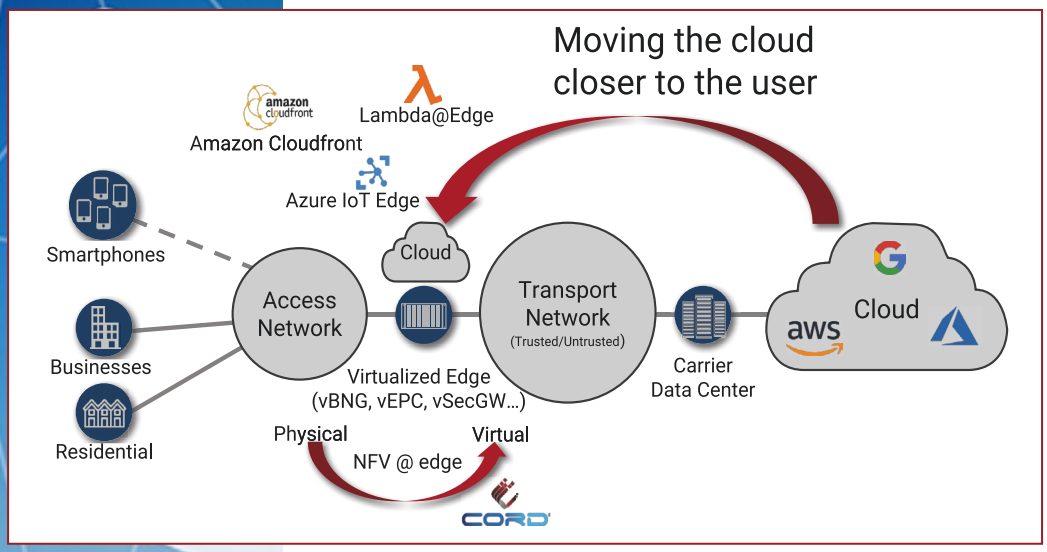

Figure 2:
The Race to the Edge

Unique Requirements at the Network Edge

There are some fundamental needs at the network edge that do not generally apply in large data centers. The most basic of these requirements is the need to reduce power and physical space, which are much more scarce at the edge than in a typical data center.

Edge sites were strategically placed far from the central network and close to the end users, in part as a cost-saving measure. These pods were not designed to host endless racks of servers, instead offering very limited physical space and a fixed power envelope. With the advent of video on top of the legacy network, the network now requires much more data, and therefore many more servers, than the existing infrastructure can handle.

Today, the number of users accessing network services at the edge is growing, and each subscriber may have multiple devices using apps that require high bandwidth and low latency. Moreover, with the growth of IoT, there can be tens of thousands of sensors sending and receiving data at any given moment, placing great strain on the network in terms of both performance and power demands.

Another major challenge at the edge is securing cyber access to the edge sites. As soon as applications are placed at the edge of the network instead of within the secure data center, the security challenges increase. Not only must applications themselves be secured, but the edge sites must also be properly isolated and user devices segregated within the network to prevent cyber attacks.

Virtual functions already tend to increase latency, even before they are merged into a single mini-data center.

As mentioned earlier, openness is also important at the edge of the network, as carriers seek to step away from vendor lock-in in their equipment. By using COTS servers that enable open implementation of services, carriers save money and ensure that their network edge remains flexible. Similarly, as new applications and new networking demands enter the market, existing hardware must be able to handle the innovations without requiring constant upgrades. When there are thousands of edge points-of-presence, changing equipment at every site is especially cumbersome and expensive in comparison to a future-proof programmable solution that can be upgraded remotely.

It is important to note that there will be multiple networking applications running simultaneously within a single edge site, and there can be applications from both carriers and cloud providers competing for precious space at the edge. vEPC, vCPE, vCDN, vRouter, vFirewall, SecGW, SD-WAN, and vBRAS/vBNG are just a few of the popular networking applications that might be found in an edge site. Virtual functions already tend to increase latency, even before they are merged into a single mini-data center.

Furthermore, as it strives to remain agile, an edge site may incorporate diverse hardware interface options and multiple networking protocols. Providing deterministic low latency is challenging when the network edge must accommodate 1G, 10G, and 25G interfaces while also handling various interconnect protocols, such as xDSL, PON, CPRI, G.fast, and DOCSIS.

SmartNICs

One solution that can suitably address the challenges of the edge is the use of Smart Network Interface Cards (SmartNICs). Legacy NICs are plugged into servers to enable connectivity, but SmartNICs add two primary features: the ability to offload data processing from CPUs and programmability of the hardware.

SmartNICs use hardware to accelerate the virtualized networking functions at the network edge and produce efficiencies that overcome the challenges to the edge sites.

Since virtual functions are traditionally run in the CPU, which adds unwanted latency, the simplest and most effective use of a SmartNIC is to offload networking workloads and functions from the CPU to the NIC. This not only reduces latency and power requirements, but also the overhead on the CPUs, freeing them for other tasks. Thus, SmartNICs use hardware to accelerate the virtualized networking functions at the network edge and produce efficiencies that overcome the challenges to the edge sites.

The market for SmartNICs is growing rapidly. According to analysis by Ethernity Networks, the overall NIC market for speeds of 10G and above is estimated at $1.4 billion in 2018, with SmartNICs accounting for 10% of sales, mostly to tier 1 data centers. However, by 2021, the overall 10G+ NIC market is expected to grow to $2.2 billion with as much as 27% of that to be SmartNIC sales. Furthermore, the shift to SmartNICs will coincide with demand for SmartNICs at the edge, with nearly half the SmartNICs being applied to telco/cloud edge installations.

There are a few types of SmartNICs:

- Multicore SmartNICs essentially add an array of processors to the card, such that networking functions are handled on the card instead of the motherboard. However, because those processors are, in essence, added CPUs on the NIC, the gains in agility, power consumption, and latency are negligible.
  Ultimately, multicore SmartNICs are ASIC-based, meaning they still rely on a proprietary element that depends heavily on the hardware vendor's delivery of new ASIC versions. In other words, a multicore ASIC-based NIC does nothing to solve the vendor lock-in issue.
  Moreover, this option is becoming increasingly less popular, as major manufacturers have begun discontinuing their flagship multicore products, such as Marvel discontinuing the Cavium LiquidIO.

- Network Processing Units (NPUs) and Graphics Processing Units (GPUs) are highly optimized to handle network functions, but they are difficult to program, tend to use proprietary microprocessors and development toolchains, and have a hard time scaling easily. In fact, most of these products have been discontinued over the past few years, including NPU-based SmartNICs from Mellanox, Microsemi, Marvell (Xelerated and Xpliant), and LSI. That trend is expected to continue.

Microsoft Azure strongly recommended against using multicore SoC NICs, and highly favored the use of FPGA-based SmartNICs.

- Field Programmable Gate Array (FPGA)-based SmartNICs are fully programmable, non-proprietary, and offer efficient, deterministic performance and scalability. They are also affordable and rapidly growing as the industry's favorite option for SmartNICs. As the market converges around this option, FPGA SmartNICs will very shortly become the de facto standard. An FPGA SmartNIC solution can also be delivered either as a pure FPGA SmartNIC or as a hybrid NIC that utilizes a standard NIC device to ease SmartNIC certification by using standard, certified software drivers.

In fact, earlier this year, in a Microsoft paper titled "Azure Accelerated Networking: Smart NICs in the Public Cloud", the company strongly recommended against using multicore system-on-chip (SoC) NICs, and highly favored the use of FPGA-based SmartNICs.[3]

3 "Azure Accelerated Networking: Smart NICs in the Public Cloud", Microsoft 2018: https://www.microsoft.com/en-us/research/uploads/prod/2018/03/Azure_SmartNIC_NSDI_2018.pdf

According to the paper, multicore NICs offer the required programmability, but come at a great cost. Latency, power, and price are limiting factors for multicore SmartNICs, and these considerations all rise precipitously when they scale beyond 40G, such that the solution is neither scalable nor futureproof.

FPGAs, on the other hand, offer performance like an ASIC and programmability like a software solution. They offer the required low latency, low power, and low price that the market demands, and they continue to do so at scale.

FPGA SmartNICs for the Virtualized Edge

Now that we have introduced SmartNICs and established that FPGAs are the superior option when choosing a SmartNIC deployment, let's revisit the unique challenges of the virtualized edge to see how FPGA-based SmartNICs address those concerns.

First and foremost, to overcome the tradeoff between space and power on the one hand and the need to accommodate thousands of devices on the other, FPGA-based SmartNICs are an ideal solution. SmartNICs can slot into existing servers, reducing the need for additional boxes and saving space. In fact, multiple NICs can reside in a single server for ultimate space efficiency. This also significantly cuts down on power requirements, as FPGAs are extremely power-efficient.

Moreover, the scalability offered by such FPGA SmartNICs enables service providers to more easily handle the large numbers of subscribers and devices at cost. Again, as Microsoft pointed out, this is only true of FPGAs, and not of multicore SmartNICs, which cannot handle the scale without significantly adding latency and power.

FPGA-based SmartNICs also address the security requirement of the edge, enabling network isolation and user segregation to prevent attacks on edge sites and user devices, as well as offering IPSec VPNs and tunnel termination.

FPGAs are open, programmable hardware. There are two primary vendors (Intel and Xilinx), and Microsoft specifically commented on how easy it was for them to port their code from one vendor's FPGA to the other and between FPGA generations. FPGA SmartNICs are a perfect complement to COTS servers, in that they are general purpose and agile. Their full programmability at the speed of software development futureproofs edge sites such that hardware does not need to be replaced or upgraded as frequently.

Finally, and perhaps most importantly, FPGAs provide highly deterministic performance with especially low latency. FPGAs can be programmed to handle all interfaces and protocols without affecting latency at all.

Comparison: FPGA-Accelerated vs. Software-Only vBNG

As an example of how an FPGA-based SmartNIC can benefit a virtualized edge site, let's compare a virtual Broadband Network Gateway application when it is run as a software-only implementation and when it runs on an FPGA SmartNIC.

An entry-level software-only vBNG deployment requires an expensive 24-core server and a large investment in licenses. This enables about 8,000 subscribers with a modest average of 3Mb per subscriber (per direction) resulting in 25G for user traffic. With the aforementioned explosion in data usage, this model is likely to be insufficient within even a few years. With a jump to only 10Mbps, the deployment would only support 2,500 users.

Compare that to an FPGA-accelerated solution. The FPGA SmartNIC can be installed in a much less expensive 8-core COTS server, which handles the data plane for transport. Only 1 of those 8 cores is actually required for the same 50G (25G in each direction) of throughput, and the same single core can scale to support 100G as well. This entry level solution will cost less than 60% of the software-only solution and only about 25% of the power[4].

Furthermore, according to a recent study, a typical mobile gateway virtual network function (VNF) produces between 129 and 1,474 microseconds of latency[5], and the jitter depends heavily on the load. Ethernity's new FPGA-based ACE-NIC100 offers twice the throughput (100G, that is, 50G in each direction) and only 15 microseconds of deterministic latency (no jitter).

Moreover, the difference is magnified when more advanced deployments (100G and 200G) are considered, as detailed in the following graphic:

Figure 3: FPGA-Accelerated vBNG vs. SW-Only vBNG

This shows that even for an entry-level deployment of 25G of user traffic, it makes sense to initiate the deployment with an accelerated VNF solution, as it is superior to a software-only solution economically, in power consumption, and in its ability to scale and plan for the future.

4 Comparative data is based on Ethernity Networks' internal lab tests in relation to data from "Finding an Efficient Virtual Network Function Architecture for Next-Generation Telecommunications Infrastructure" by Intel, Hewlett Packard Enterprise, and China Telecom's Beijing Research Institute: https://www.intel.com/content/dam/altera-www/global/en_US/pdfs/literature/wp/wp-01273-finding-an-efficient-virtual-network-function-architecture.pdf

5 "Validating Nokia's IP Routing & Mobile Gateway VNFs", Light Reading and the European Advanced Networking Test Center AG (EANTC): https://www.lightreading.com/ethernet-ip/new-ip/validating-nokias-ip-routing-and-mobile-gateway-vnfs/d/d-id/720902

Multi-Access Edge Computing Option

The previous section showed how an FPGA-accelerated VNF uses dramatically fewer cores than a software-only solution. This is significant when an edge site is running multiple VNFs simultaneously. Ethernity Networks offers a solution in which multiple FPGA SmartNICs are installed in a single server, which can serve many virtual networking applications at once, all within the same data plane. This provides the ultimate in space, power, and latency efficiency.

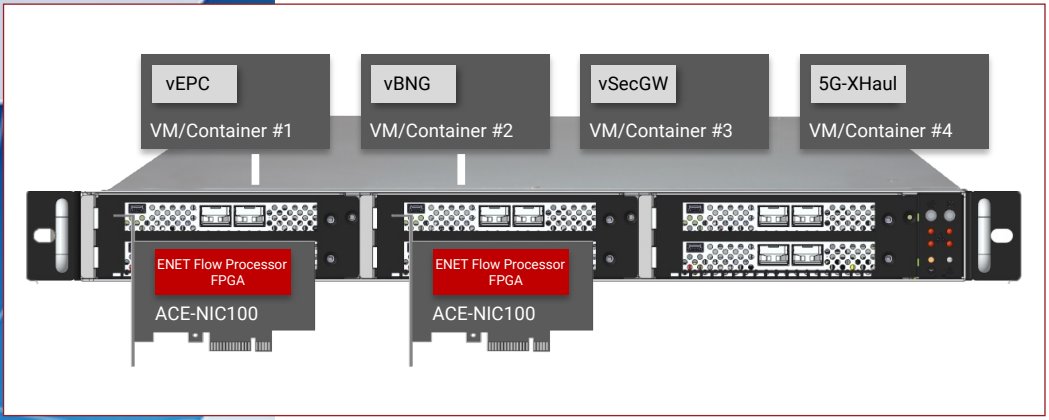

Figure 4:
Ethernity's Multi-Access Edge Computing Solution

Conclusion

As communications service providers and large central cloud providers extend their virtual infrastructure toward the edge, technology is required to overcome the unique challenges that arise. When virtual networking applications are run through software programmed on legacy hardware, service providers face severe challenges of space and power to reach the growing number of users and devices, and they are limited by performance issues, especially in terms of latency and jitter.

An FPGA-based SmartNIC solution offers scalable, deterministic performance with very low latency at a fraction of the space, power, and overall cost.

The best way to address these concerns is to implement an FPGA-based SmartNIC solution, which offers scalable, deterministic performance with very low latency at a fraction of the space, power, and overall cost. Such a solution is secure, open, and fully programmable, ensuring that costs remain low well into the future.

About Ethernity Networks

Ethernity Networks (AIM: ENET.L) is a leading innovator of network processing technology and products. Mounted on low-cost COTS FPGAs and with a rich set of networking features, Ethernity's ACE-NIC smart network adapters, ENET SoCs, and network appliances offer best-in-class all-programmable platforms for the fixed and mobile telecom, enterprise security, and data center markets. Our complete offering, incorporating hardware, FPGA firmware, and software applications, enables full programmability at the pace of software development, quickly adapting to changing market demands and applications and facilitating the deployment of edge computing, 5G, and SDN/NFV.

ETHERNITY NETWORKS

For more information please visit us at www.ethernitynet.com