FPGA as a Platform for Complete Network Disaggregation

ETHERNITY NETWORKS

September 2020

CONTENTS

1 NFV as a Precedent 3
2 Disaggregation Trend 4
3 The Next Frontier: Network Disaggregation at the Edge 4
4 The Winds of Change 6
5 The FPGA as a De Facto Platform 7
6 Use Case – Microsoft Azure 9
7 Use Case – Open UPF 10
8 Costs vs. Benefits 12
9 Summary 12

1 NFV AS A PRECEDENT

The concept of Network Function Virtualization (NFV) has advanced beyond the theoretical to enable enterprises to eliminate the costs associated with procurement and maintenance of hardware equipment in enterprise and small and medium business (SMB) usage. Enterprises are rapidly embracing network virtualization solutions, as they promote faster server provisioning and quick deployment of network services. SMBs are also expecting heavy usage of virtualization solutions to enable consolidation of servers and applications and improve disaster recovery. By deploying these solutions, they can also handle multiple workloads with maximum uptime and improved performance.

NFV offers much greater agility in addressing evolving requirements, provides added control over both current and future functionality, and eliminates vendor lock-in. It can therefore generate significant savings in an enterprise's operational expenses.

Rapid investments in the commercialization of 5G networks will also drive the demand for NFV solutions. The implementation of virtualization technology will help to enhance a 5G network's functional and architectural viability, including increased agility and reduced capital expenditure.

However, there are a few major obstacles to NFV deployment in the telco/edge cloud – especially for 5G – including the complicated move from rigid network equipment to a fully virtualized environment. Network performance when relying on NFV tends to suffer, especially through added latency compared to hardware-based networking.

Nonetheless, the market has determined that the benefits outweigh both the performance issues and the added capital expenses. The flexibility that NFV provides, futureproofing the network from the need to replace field-deployed hardware and enabling choice of vendors, is well worth the initial investment.

This sets a precedent for other areas of the network. If it makes sense to pay more up-front to avoid vendor lock-in and gain agility and control when it comes to NFV, why shouldn't that same model be used when it comes to white box edge switches and routers?

2 DISAGGREGATION TREND

NFV is part of a larger disaggregation trend that has enabled the industry to move toward agile networks. Traditionally networks relied on ASIC (application-specific integrated circuit)-based monolithic hardware appliances that bundled proprietary software into a vendor-locked device and integrated only with other offerings from that vendor. Disaggregation has changed the paradigm to overcome such limitations.

Thanks to the use of X86 as a standard platform for server hardware, today's networks can disaggregate the software applications from the underlying hardware bare metal server. Software-based functions run on top of CPUs inside standard servers from any of several different vendors, and open stacks that communicate between virtual machines for application and service chaining dictate the overall appliance functionality that runs on the server.

Disaggregation within the data center provides the flexibility to choose a software vendor separately from the choice of a hardware server platform. It ensures that the resulting appliance is futureproof because the software can be upgraded or replaced without needing to replace the hardware that hosts it.

3 THE NEXT FRONTIER: NETWORK DISAGGREGATION AT THE EDGE

While disaggregation has become standard in data centers and cloud networks, it has yet to penetrate fully at the network edge. There has been a separation of hardware and software vendors, leading to the usage of white box edge switches, but the complete revolution in disaggregating the hardware has yet to occur.

Operators have gained some measure of choice in their software vendors and greater configurability of the control plane, although data plane functionality within carrier-grade white box switching and routing hardware is still limited by the ASIC-based silicon. Even supposedly open platforms that offer a choice of operating system or control software still maintain the vendor lock-in when it comes to the underlying hardware-implemented data plane. Moreover, that data plane functionality often lacks many of the advanced features that are required by operators at the network edge.

Worse yet, the network edge is exactly where disaggregation is most needed. With the standards and requirements of the edge still evolving, it is impossible to know today what will be required even a year from now, so perhaps the most important characteristics for operators to consider in selecting edge switching hardware should be agility and futureproofing. The one certainty at the edge is that the requirements will change, and yet, ASIC-based switch hardware, even in white box devices, cannot be upgraded to address new data plane functionality requirements.

Adding new hardware to account for a lack of functionality is not always an option at the network edge, where both physical space and power are at a premium. The alternative – replacing field-deployed hardware – is prohibitively expensive, costing up to five times the initial capital expense. Even if it were affordable to do so, replacing ASIC-based hardware with new or upgraded ASIC-based hardware guarantees nothing in terms of futureproofing. It could be that in another year, the operator will need to upgrade yet again.

The 2020 State of the Edge report confirmed the existing approach being taken by operators:

"The edge computing market and its infrastructure demands today are not the same as they will be in the future. Today, edge computing does not support a 'one-size-fits-all' approach to infrastructure investment. Edge platforms are distributed with scarce resources relative to the seemingly infinite resources of hyperscale centralized cloud data centers. Until edge computing matures with large scale implementations, deployments will be use case dependent, often with proprietary hardware and software platforms, because of resource scarcity. This is in stark contrast to hyperscale data centers that are built on massive general-purpose platforms with the scale needed to support virtually any use-case." [1]

Perhaps the issue is that no standard platform has emerged to enable full disaggregation of the data plane hardware. Whereas in the first-level disaggregation occurring in data centers and cloud networks there is an X86 standard upon which any vendor can provide disaggregated software, no such standard has yet become the de facto generic platform for data plane functionality and feature development. The market remains dominated by a few large vendors who maintain a strong relationship with the original equipment manufacturer (OEM) channels.

[1] State of the Edge 2020, Jacob Smith and Matt Trifiro, editors, p.14

4 THE WINDS OF CHANGE

Despite the position taken by State of the Edge that resource scarcity is driving operators to rely on proprietary hardware at the edge, industry purchasing behavior has begun to change. Recent bans on Huawei equipment in 5G networks in the United States and the United Kingdom have forced operators to consider alternatives to their go-to inexpensive ASIC provider. These governmental actions, in conjunction with the emergence of independent software vendors that are programming specifically for edge applications, create an excellent opportunity for operators to re-evaluate their networks and consider alternatives.

5G infrastructure is a green field, with virtually no carry-over from legacy equipment. As such, operators can now consider how to meet the insatiable demand for high bandwidth and low latency for 5G, Internet of Things (IoT), and edge computing with an eye to the future and less concern about the expense involved.

The solutions currently dominating the marketplace fall short in a number of ways. While the operator can select a software stack that supports all the required networking features, the underlying ASIC-based switch silicon could be a limiting factor. While most switches have a decent forwarding scale, many advanced features may be lacking. It is important to ask some basic questions before committing to a hardware purchase:

- Can new functionalities be added easily?
- Can the switch be migrated for IoT aggregation?
- Are all new tunneling protocols supported?
- How are new security algorithms or security holes handled?

But the most important question has become:

- Why are we still using ASIC-based hardware for switching/routing at the edge?

Perhaps the current answer is that there is a lower capital expense incurred by using ASICs. That would account for the perplexing decision by multiple UK operators to replace Huawei equipment with similar ASIC-based equipment from Ericsson or Nokia. But considering that there are issues with rigid functionality and vendor lock-in with any ASIC-based solution, why are operators not using this golden opportunity to disaggregate and bring out the full potential of 5G?

In fact, at the edge it is even more important to achieve maximum flexibility in both the control and data planes, and, considering the evolving nature of the requirements in edge applications, futureproofing is exceptionally valuable. The long-term savings in operating expenses by fully disaggregating the hardware will far outstrip the added capital expense involved in migrating away from ASICs.

The time has come for truly disaggregated networking solutions at the edge.

5 The FPGA as a De Facto Platform

One technology already being used to a large degree for accelerating the data plane in the NFV market is the Field Programmable Gate Array (FPGA). FPGAs are programmable hardware optimized to handle networking and security traffic and which can be incorporated into a white box edge device by way of a system-on-chip (SoC) or smart network interface card (SmartNIC).

Many networking and security functions are highly CPU-intensive, burning through CPU cores with varying levels of performance. When the data functions are offloaded from CPUs to FPGAs, performance not only improves; it also stabilizes and is more deterministic, and it enables the CPU to serve as a co-processor, saving cores for the compute and control functions they were meant to handle.

Furthermore, FPGAs offer complete network hardware disaggregation. Operators can select the required FPGA just as they would select the required CPU and processing power in a server. FPGAs are available from multiple vendors, all of which offer development toolkits. Operators would also have the freedom to choose the code that runs on an FPGA-based network interface card, selecting from an array of firmware vendors. It is also easy to port IP from one FPGA to another.

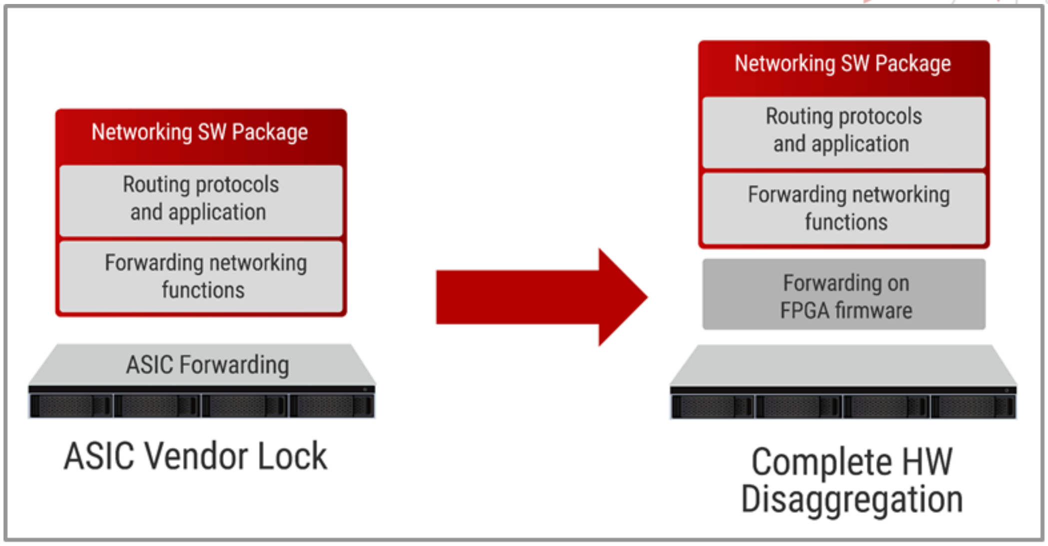

Figure 1: The FPGA is an ideal platform for handling data forwarding while avoiding vendor lock and the lack of flexibility of ASICs

FPGAs maintain the philosophy of disaggregation that was started with NFV. While FPGAs currently cost slightly more up-front in capital expenditure than ASIC-based devices (at least until 7nm FPGAs become the standard), they provide a huge advantage in that they add agility and futureproofing to the network edge. This is crucially important – at the very least until the standards and benchmarks required at the edge are formalized – and probably well beyond that.

FPGA-based SmartNICs offer these significant advantages over ASIC-based hardware:

- **Flexible** – FPGAs can be used to address a wide range of acceleration use cases, including compute, networking, and storage.
- **Adaptable** – FPGA firmware can be redesigned very quickly, and new functionality can be added on top of the existing IP. When compared to the years it often takes for an ASIC design cycle, operators can save time and money by using FPGAs instead of waiting for the fabrication of each new ASIC.
- **Programmable and Reconfigurable** – The same firmware can be repurposed to handle new applications, or the mix of workloads on a single piece of hardware can be reconfigured as needed as conditions change. There is rarely a need to replace or add hardware.
- **Performance** – FPGAs are deeply pipelined, with massive on-chip memory and external memory to DRAM. They can therefore deterministically provide high throughput with extremely low latency and jitter.

As a result, FPGAs are an ideal platform for truly disaggregating hardware at the network edge, continuing the trend that was begun with NFV. They extend the concept to apply to white box edge switches and routers as well as to optical line terminals (OLTs). FPGAs perfectly address the concern about the use of proprietary ASIC-based hardware platforms while also improving performance, avoiding vendor lock-in, and futureproofing the network, thereby saving on long-term operating expenses and reducing total cost of ownership.

6 Use Case – Microsoft Azure

When it comes to innovating using FPGAs, no one has greater expertise than Microsoft Azure. As early as 2010, well before anyone had ever heard of a SmartNIC (a term Microsoft Azure actually coined), Microsoft demonstrated the first proof of concept that used FPGAs at scale to accelerate web searches on its Bing search engine. The company's Project Catapult produced the first pilot program to deploy FPGA-enabled servers in a data center in 2013, which showed dramatic improvement in latency, running decision-tree algorithms 40 times faster than CPUs alone while reducing the number of servers.

By 2015, Microsoft deployed FPGAs at scale into its Azure public cloud, and within a year, its AccelNet program had introduced FPGA-based SmartNICs as the default hardware for implementing virtual network functions in Azure, deploying FPGAs in over one million hosts.

In 2017, Microsoft released a white paper on Accelerated Networking that made a compelling case for FPGA-based SmartNICs as the optimal choice for cloud environments. Many of the points made therein are relevant to the network edge as well, including:

1. While ASIC-based devices offer the highest performance potential, they lack programmability, which limits their adaptability over time. Adding a small number of CPU cores to an ASIC does not provide the required programmability, as those CPU cores quickly become the bottleneck.
2. A strong recommendation against using multicore system-on-chip (SoC) NICs, which 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 the network scales beyond 40G, such that the solution is neither scalable nor futureproof.

3. Adding additional host cores to manage pure software solutions offers the most flexibility in terms of programming, but sacrificing host cores to run the virtual switch, even with technologies such as DPDK, is too costly, both in terms of dollars and in terms of performance. Burning those cores essentially kills the business model.
4. FPGA firmware based on pipeline architecture, on the other hand, offers the performance and efficiency of a customized ASIC while enabling the programmability to adapt to new features, the ideal balance for optimizing the network.

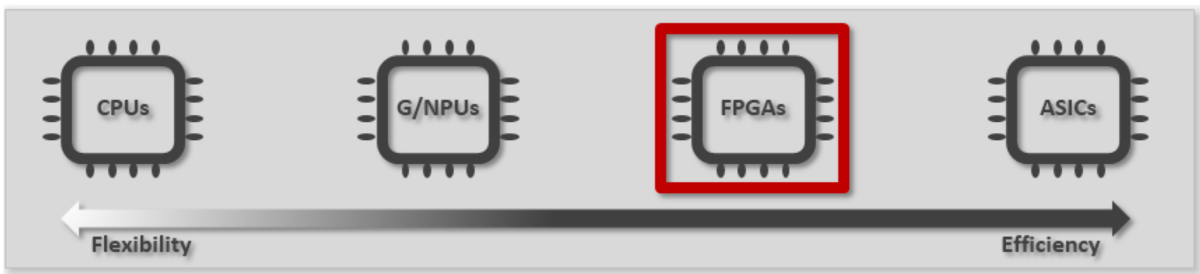

Figure 2: Microsoft Azure maintains that FPGAs provide the optimal combination of performance and flexibility

All of Microsoft's messages about FPGAs are amplified at the network edge, where the benchmarks and standards have yet to be finalized, necessitating an agile, futureproof solution. The usage of FPGAs as a de facto platform for edge switch/routing is an ideal way of overcoming the uncertainty of a burgeoning market while ensuring high performance and maximum flexibility.

7 Use Case – Open UPF

In January 2020, China Mobile gathered the leading 5G hardware and software vendors to an "Open User Plane Function" conference, where they addressed the need to create a universal platform for user plane functionality (UPF), which is the data plane within a 5G network.

This signaled that the world's largest telecommunications operator believes that UPF appliances need to be both open and decoupled from the rest of the 5G infrastructure to increase the flexibility of the network infrastructure and facilitate operation at the network edge.

China Mobile took this initiative because it recognizes that UPF can make or break a successful rollout of 5G. As 5G demands higher bandwidth, lower latency, and other performance improvements over 4G LTE networks, operators like China Mobile have decided that the user plane, where the data itself passes through, must be moved to the edge. Control and configuration, on the other hand, can be kept in a central location. In order to make this work, China Mobile is asking vendors to work together so the UPF can be open and separated from the software control plane.

China Mobile further believes that FPGA firmware is the right programmable platform to handle the user data plane. They called for delivery of an FPGA SmartNIC to offload the data plane from the server CPU, further speeding up and improving the performance of the network. When it comes to creating a universal platform to handle the user plane function, China Mobile believes this can be accomplished best with an FPGA, which provides the following advantages:

- Far better performance (throughput, latency) than standard CPUs, which are simply not built for network data flow processing
- Small footprint and low power consumption, ideal for the network edge
- Runs on standard off-the-shelf servers, where the embedded FPGA itself is a generic platform, as opposed to the dedicated hardware of an ASIC
- Reprogrammability as needed to meet evolving standards and requirements

China Mobile's initiative is an excellent example of the trend toward complete hardware disaggregation in the network, separating control and user functions for true network function virtualization, while keeping the infrastructure agile enough to adapt to ever-changing standards and requirements.

China Mobile's Open UPF policy will also allow more Tier 2 vendors to compete to supply 5G solutions, as the operators will be empowered to select the optimal solution for UPF rather than rely on the Tier 1 ASIC suppliers for entire systems.

A similar and better publicized initiative is O-RAN (Open Radio Access Network), which seeks to disaggregate the primary edge of the 5G network. FPGA SmartNICs are equally as effective as a platform to enable open, high-performance data processing throughout the O-RAN, including the Centralized Unit and the Distributed Unit, as they are in handling UPF.

8 Costs vs. Benefits

As mentioned, FPGAs currently require a slightly higher capital investment than ASIC-based edge hardware devices. However, they can save operators significantly on operating expenses by eliminating the need to replace existing hardware with upgrades or with new devices that include a broader feature set. Replacing field-deployed hardware to support new functionality can incur costs of up to five times the initial cost of the hardware itself.

However, it is important to note that just as CPUs improve with each new generation and enable higher software performance for the same price, FPGAs continue to improve with each new generation as well. This increases the amount of logic available in FPGAs over time for the same price, further improving performance without adding cost over previous generations.

Moreover, FPGA prices are dropping steadily, due to the scale of production as FPGAs gain in demand and capture more of the traditional ASIC market, and due to advanced silicon node production. The gap between the capital expense of FPGAs and ASICs is already greatly reduced and will continue to lessen with 7nm FPGAs. This is even more true when comparing logic capacity to price.

9 Summary

NFV has ushered in a trend of disaggregation in networking hardware in the data center and cloud despite higher capital costs in order to achieve agile, open platforms that save network administrators on operating expenses in the long run. Similarly, telecommunications operators should now be seeking further disaggregation of their switch/routing hardware at the network edge, where white box edge switch/routers and OLTs currently still rely on rigid ASIC-based switch silicon from a clique of vendors who act as if they own a monopoly on the switching marketplace.

FPGAs offer ASIC-like performance with the programmability of software. They are increasingly affordable, and they are rapidly improving their capacity/price with each new generation.

As with the philosophy behind NFV, complete network disaggregation sacrifices a bit of up-front cost in exchange for the benefit of significantly lower TCO in the long run. By establishing FPGAs as the optimal open standard platform for handling the data plane in edge devices, operators gain tremendous flexibility to overcome vendor lock-in and the programmability to futureproof their networks, the ultimate requirement for today's multi-access edge compute network.

About Ethernity Networks

Ethernity Networks Ltd. (AIM: ENET) provides innovative, comprehensive networking and security solutions on programmable hardware for accelerating telco/cloud networks performance. Ethernity's FPGA logic offers complete Carrier Ethernet Switch Router data plane processing and control software with a rich set of networking features, robust security, and a wide range of virtual function accelerations to optimize telecommunications networks. Ethernity's complete solutions quickly adapt to customers' changing needs, improving time-to-market and facilitating the deployment of 5G and edge computing.

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