NCNF: From "Network on a Cloud" to "Network Cloud"

The Evolution of Network Virtualization

WHITE PAPER

Introduction

DriveNets Network Cloud allows communications service providers (CSPs) and cloud providers to evolve their infrastructure and enable the implementation of an efficient edge cloud architecture that also supports latency-sensitive applications. Moreover, Network Cloud delivers performance improvements and significantly reduces total cost of ownership (TCO) by using telco assets in an optimal way.

The DriveNets Network Cloud model addresses the same network, service-scaling and time-to-market challenges that network function virtualization (NFV) has tried to solve. It completes the journey from physical network functions (PNFs) and virtualized network functions (VNFs) through cloud-native network functions (CNFs) and now to Network Cloud network functions (NCNFs).

This journey started with VNFs that eliminated some limitations of PNFs, but were not geographically scalable. It continued with the introduction of CNFs that addressed geo-scalability but were not networking optimized. It now has reached ultimate optimization by “building networks like clouds” with the creation of NCNFs.

Virtualized Network Functions (VNFs)

Building networks is complex. The notion of simplifying it by disaggregating network functions from the hardware they run on is not new. Network function virtualization (NFV) introduced the concept of the virtualized network function (VNF) that enables replacing a monolithic solution with a software implementation of a network function over an x86-based server.

Decoupling PNF into software (VNF) and commodity hardware

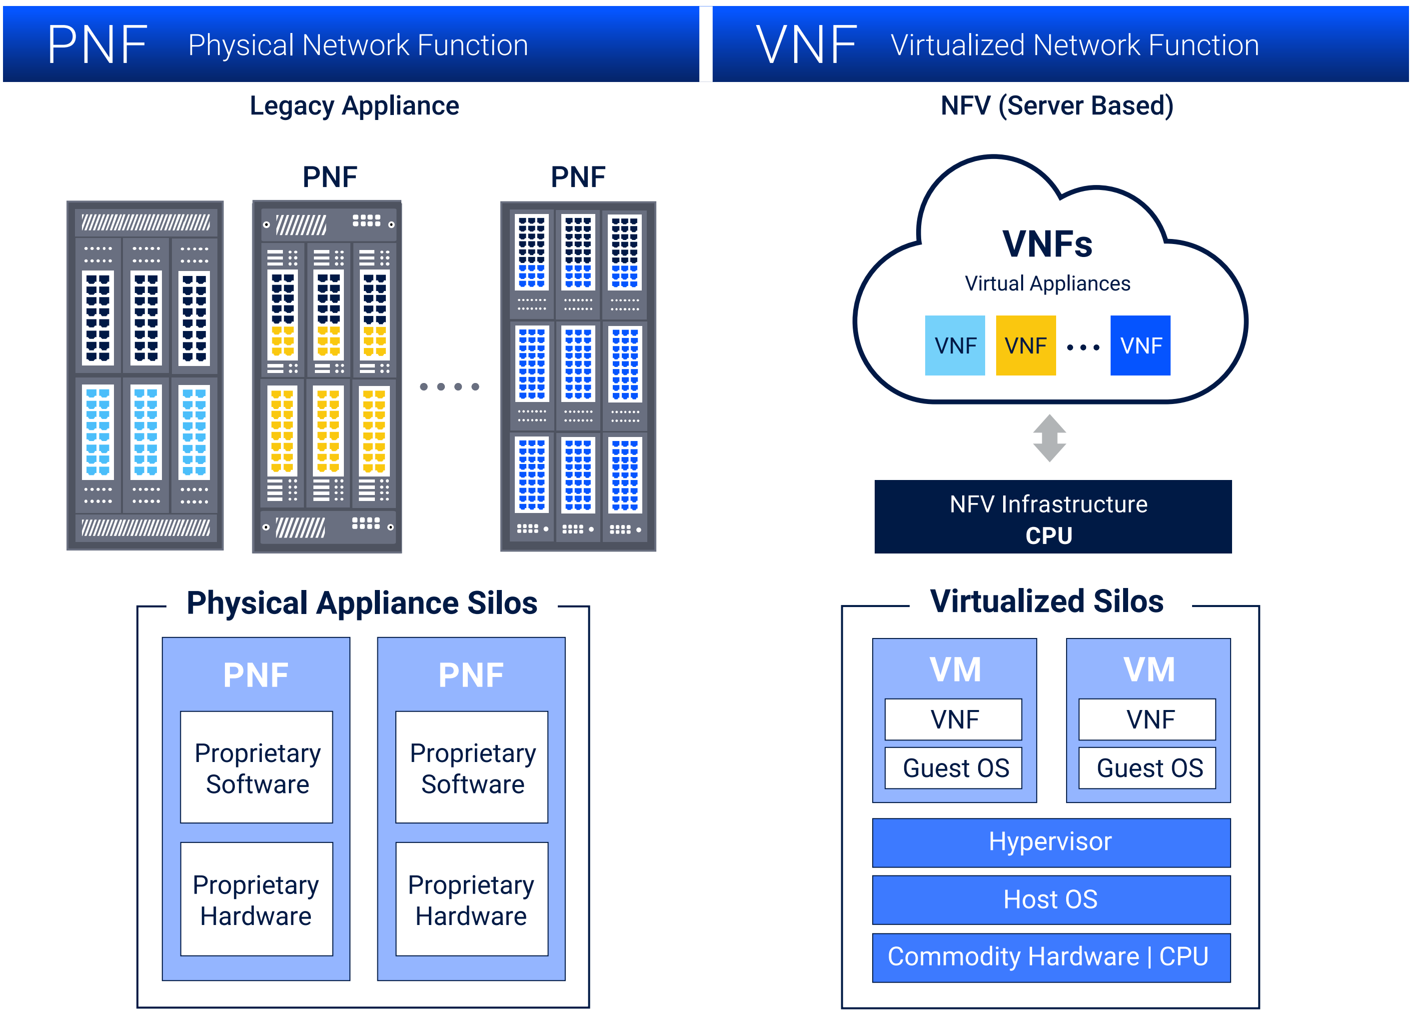

The previous physical network function (PNF) architecture is based on monolithic network elements, purpose-built to implement a single network function such as a router (e.g. core, edge or access), a firewall, or deep packet inspection (DPI). This architecture is extremely non-cost-efficient as it requires multiplication of hardware elements across the network. VNFs eliminate the limitations of the monolithic chassis by implementing the network function over a standard, scalable, unified hardware infrastructure, i.e. x86 servers.

While providing some value to service providers, VNFs lack the ability of easily deploying and scaling network functions (NFs) across multiple sites. With VNFs, adding a networking site requires significant investment in IT infrastructure, which blocks VNFs from being utilized and scaled on a global scale.

Cloud-Native Network Functions (CNFs)

The next step, therefore, was the move to cloud-native network functions (CNFs). This architecture allows an elastic implementation model for NFs, including dynamic resource allocation and even porting of these functions according to service needs. Developed using microservice architecture, CNFs can run, essentially, in containers over any cloud infrastructure. This gives operators a truly elastic network in which they can implement functions over a public, private or hybrid cloud.

Decoupling the monolithic app (VNF) into microservices (CNF)

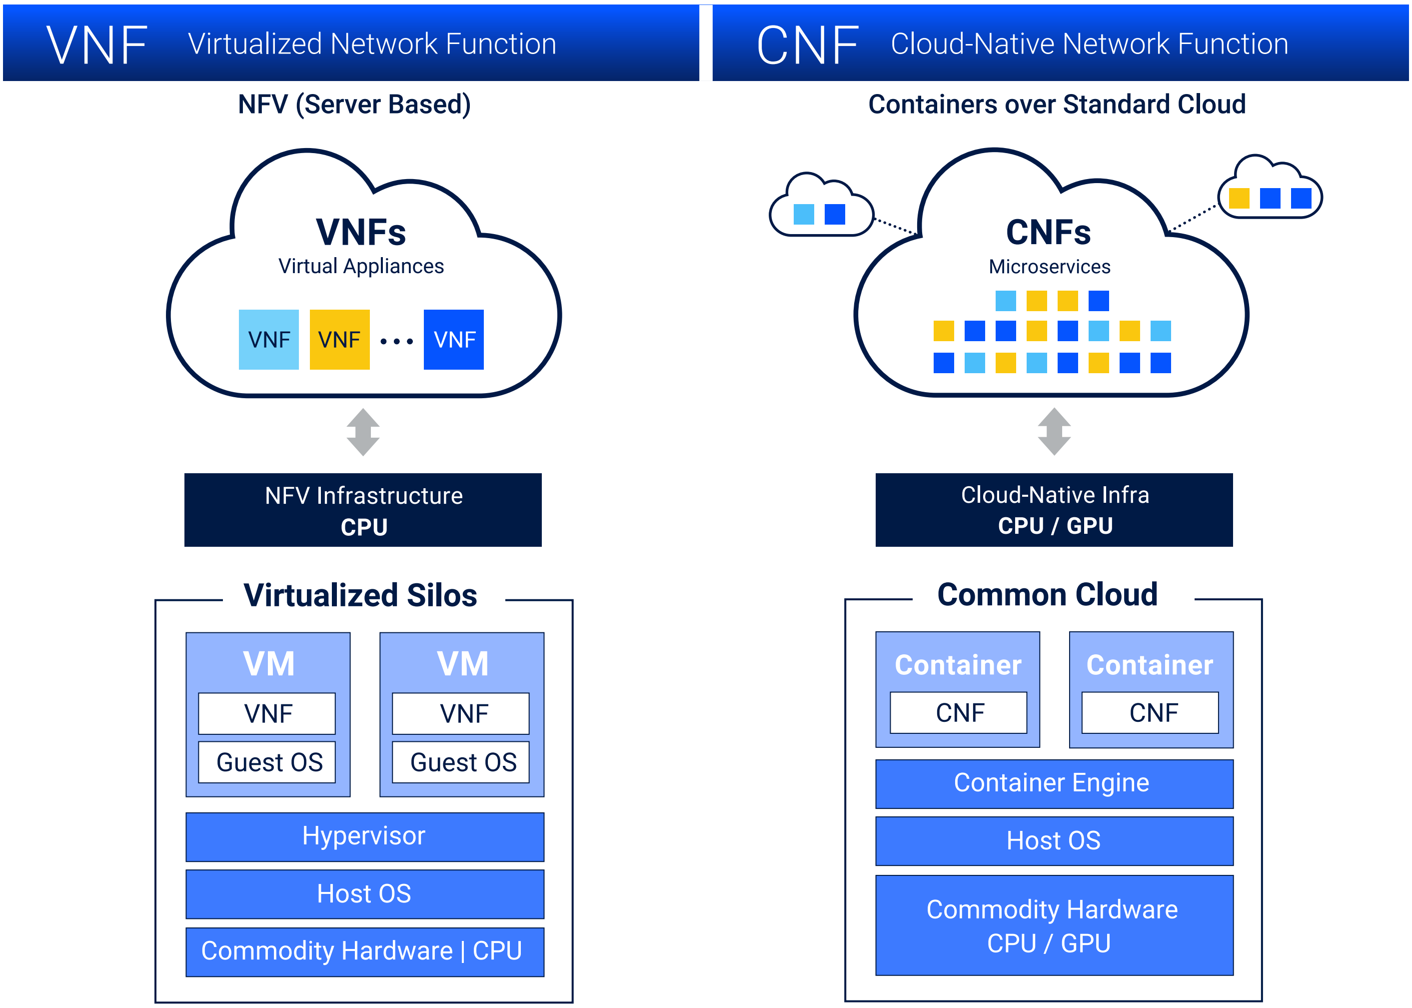

CNF Limitations

While running network functions over cloud instances increases service scale and flexibility, it also introduces great inefficiencies and high costs. A server-based cloud architecture utilizes CPU (central processing unit) and GPU (graphic processing unit) resources, which are optimized for wide-ranging applications but not for networking functions. This means that when running a CNF over a CPU/GPU infrastructure, performance falls short (e.g. in terms of forwarding capacity, queueing and buffering, and access list scalability), and service quality can be degraded. This substantially limits the introduction of new services, especially latency-sensitive ones that are classified as ultra-reliable low-latency communications (URLLC) in the 5G arena.

These limitations are extremely hard to resolve. Gathering enough processing power to perform network functionality means stacking a large number of servers; this is inefficient in terms of power and floor space, and is not feasible in the network edge where real-estate is limited.

Harnessing networking-optimized network processing units (NPUs) to perform these network functions can efficiently address scaling of network-aware services. Unfortunately, the platforms supporting them, mainly routers, are proprietary and closed, and do not offer the APIs and abstraction layer needed to run NF microservices as containers efficiently at scale.

Network Cloud Network Functions (NCNFs)

DriveNets Network Cloud enables the long-awaited solution to this scaling limitation. Network Cloud is a cloud-native architecture based on a cluster of networking white boxes, which incorporate CPUs and NPUs to address the networking-specific requirements of network functions. These NFs are implemented over the networking-optimized cloud in an architecture called Network Cloud network function, or NCNF, which utilizes a combined CPU and NPU resource pool.

Running NF on networking-optimized (NPU) cloud instances

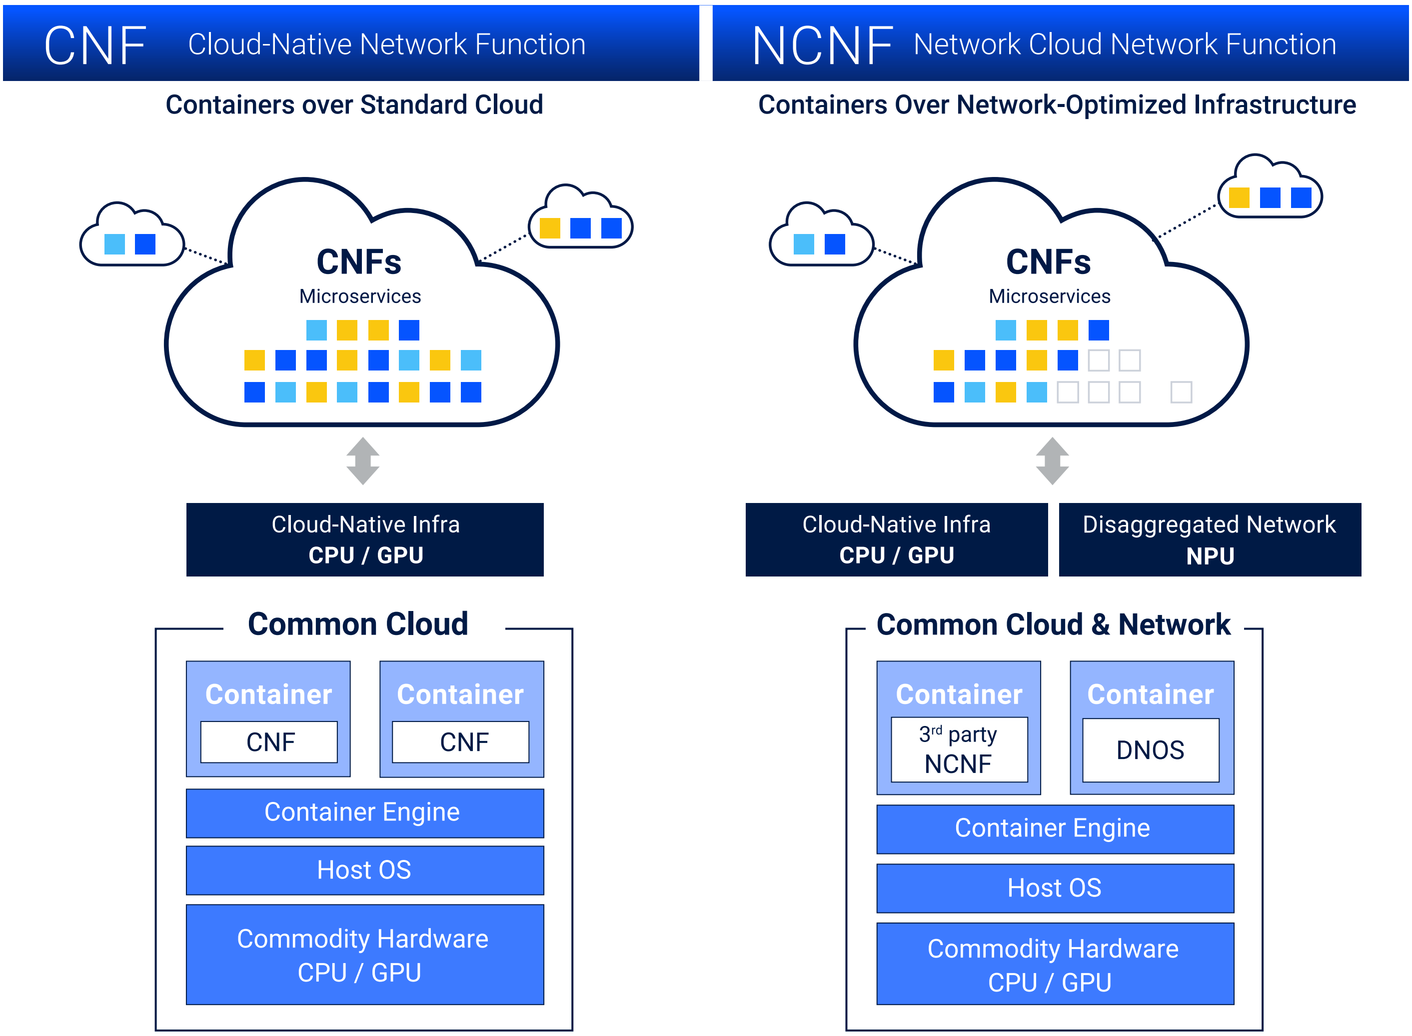

Networking-Optimized White Box Cluster

The NCNF architecture allows operators to efficiently implement functions that require intense networking resources over a lean infrastructure of CPUs and NPUs. This is achieved by sharing resources across different instances of the hardware infrastructure, thus allowing optimal utilization of bandwidth resources, x86 resources, QoS resources and TCAM resources, which are typically not balanced by a single service.

The following diagram illustrates the different resource requirements for different network services:

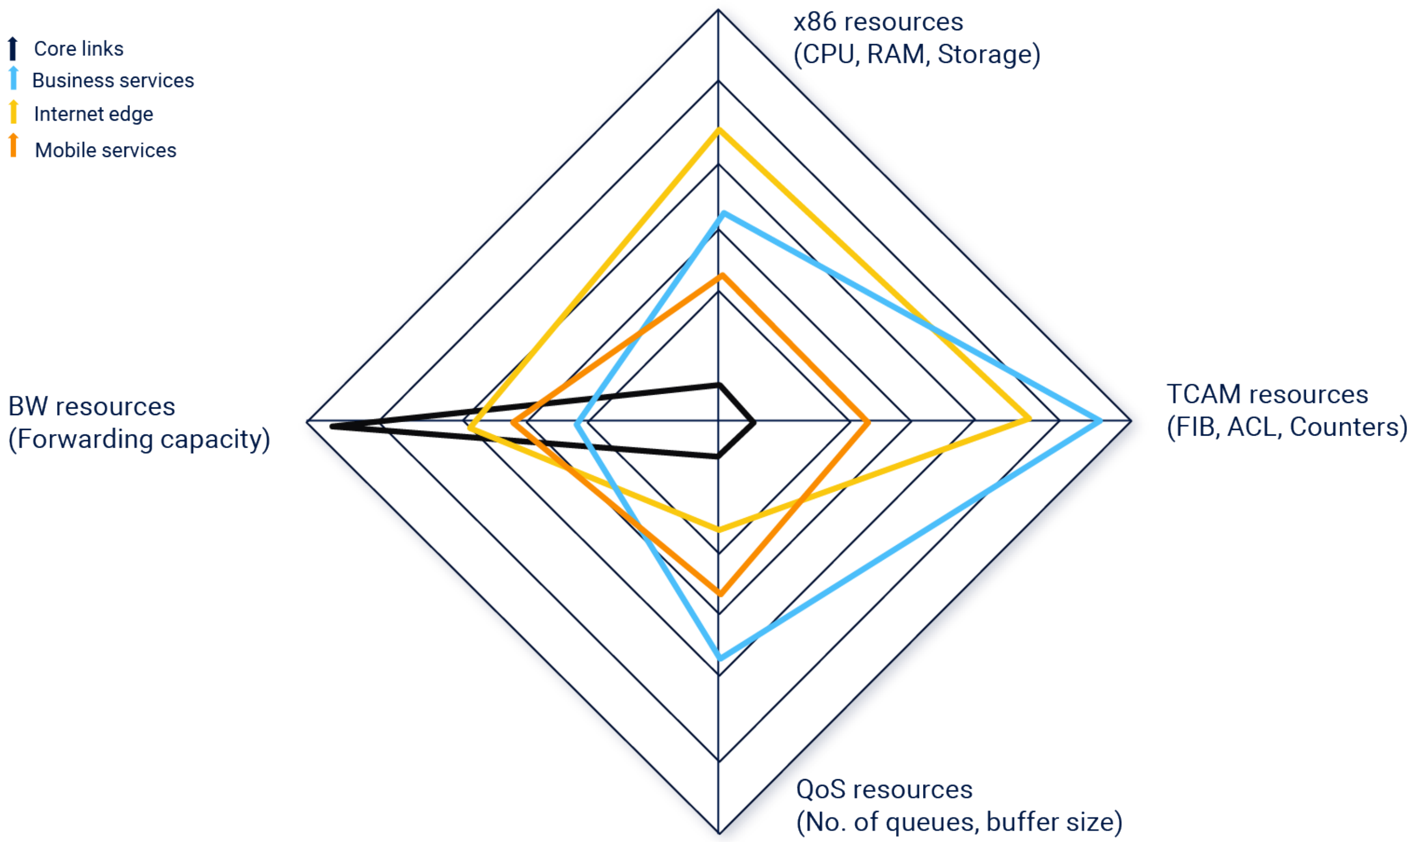

Resource sharing allows full utilization of infrastructure resources. For example, infrastructure can be shared between core and business services networking functions, which represent a highly complementary requirement map for BW and TCAM resources.

This implementation of NFs over the NCNF infrastructure not only improves resource utilization but also enhances service quality, by using optimized NPUs for networking functions. It also allows the tight integration of additional NFs with the networking infrastructure. These NFs can include firewall, DDoS prevention, DPI, content delivery network (CDN), and Evolved Packet Core/5G Core (EPC/5GC), among others.

This tight integration leads to another major benefit – the ability to dynamically adjust and optimize NFs based on a wide array of traffic and performance metrics collected, in real time, by the NCNF infrastructure. For instance, this could enable a significant reduction in remedy time and an increase in scale of a DDoS prevention platform, or change how a CDN distributes streaming content according to specific performance parameters in every network edge zone.

NCNF architecture offers the following key advantages:

- Efficient sharing and optimal utilization of network resources
- Enhanced service quality via optimized NPUs for networking functions
- Tight integration of NFs with networking infrastructure
- Dynamic optimization of NFs via real-time collection of traffic/performance metrics

The Edge Cloud – the New Edge

All of the forementioned advantages of the NCNF architecture become essential when it comes to time/latency-sensitive applications. Such applications are gaining momentum with the introduction of the 5G URLLC use case and the ongoing development of Industry 4.0 remote operations applications, as well as gaming, augmented reality (AR) and virtual reality (VR) apps.

These applications do not yet represent a major portion of traffic running through the CSP network, though this may change with the massive introduction of AR/VR devices and the maturing of level-5 autonomous vehicle technology. However, these apps represent a fast-growing portion of operators' revenues. Moreover, they require end-to-end time-sensitive networking (TSN) performance, which is changing the way the network edge is planned and built.

The new network edge is no longer only about compute and storage resources; it is, more than ever, about networking resources. As such, the advantages of NCNF architecture become a must in any edge implementation, making the concept of the "edge cloud" the networking-aware successor of the edge-computing function.

Conclusion

There is a new way of building networks for communication service providers, hyperscalers and other network operators. The DriveNets Network Cloud architecture is essential in any modern network as is it the only one that can integrate compute and networking resources while addressing service providers' cost, scaling and innovation challenges.