HEAVY READING

WHITE PAPER

Independent market research and competitive analysis of next-generation business and technology solutions for service providers and vendors

Disaggregation and the Virtual Network

A Heavy Reading white paper produced for Intel

AUTHOR: JIM HODGES, CHIEF ANALYST, HEAVY READING

INTRODUCTION

In late 2012, communications service providers' (CoSPs') network evolution strategies were at a crossroads. Not only were traditional networks facing capacity and service agility constraints, the associated operational costs continued to escalate faster than service revenue. Consequently, many CoSPs embraced the concept of network functions virtualization (NFV) during this period when it was introduced and positioned in a joint operator white paper as the most realistic path to cloudifying telco networks.

NFV is often positioned simply as an architecture change, but it is also transformative. It defines a foundational blueprint to enable the seamless adoption of multiple technologies, such as 5G, edge computing, and other distributed cloud-based architectures. In a business transformation context, virtualization provides the flexibility to support dramatically changing carrier and enterprise business models. While NFV remains both technically and business relevant today, many CoSPs have found implementation very complex. NFV implementation will require further evolution in order to fully address operational and business requirements.

This white paper documents the factors currently driving this network virtualization evolution, including examining the impact of adopting an innovative model based on software and hardware disaggregation.

THE EVOLUTION OF VIRTUALIZATION

As referenced above, the NFV journey that CoSPs embarked upon more than 6 years ago is ongoing and continues to be shaped by the business and technical realities of the virtualized cloud era. While the transition from a traditional network design to a virtualized architecture has not been without challenges, this was not unexpected given that it represents a complete change in the way products are architected, designed, built, and operated.

Despite the challenges, NFV has been successful on many fronts. For example, in a service context, virtualization is important because it enables the adoption of advanced cloud-based technologies such as 5G. Today, 5G is not only viable, it is increasingly a commercial reality that will start to reach its commercial zenith in the next 6 years. By 2025, the GSMA forecasts that the number of 5G connections will reach 1.4 billion subscribers.[*]

Virtualization also fosters the adoption of open-source software. For a significant number of CoSPs, implementation of open-source software represents a strategic imperative because it both brings them more control and expands partner ecosystems. Such implementation permits a shift from a single, integrated turnkey solution vendor solution option to multi-vendor solutions that deliver the greatest implementation flexibility and innovation at the lowest price point.

Another area that virtualization addresses is a change in the reach of services and richness of service models. Specifically, the notion of what constitutes a CoSP's service continues to expand and evolve beyond basic access and transport. This is driven by the fact that cloud adoption is happening in adjacent industries, such as transportation, healthcare, and first

[*] GSMA, *The Mobile Economy 2019*, February 25, 2019.

responder segments, which opens the door for progressive CoSPs to enter these markets on more than simply a connectivity level by upselling security and other managed services.

At the same time, telco domain service models are also evolving. The most visible change is the focus on using 5G and other edge technologies to deliver low latency services both to their mobile and enterprise subscribers, as well those of the adjacent market segments. The ability to deliver low latency massive-scale services is often cited as the key capability that will drive rapid adoption of Internet of Things (IoT) services, ultimately accelerating the digital transformation of businesses across numerous industries and service provider landscapes.

Changes in the competitive landscape are also having a significant impact. The continued rise of over-the-top (OTT) service providers has flattened or eroded CoSP revenue, providing them with the incentive to push the services innovation curve to achieve parity with OTT service innovation performance metrics and to deliver more value to customers.

CURRENT NFV IMPLEMENTATION TRENDS

Given the impact of these factors, it is not surprising that 2019 marks the beginning of a new NFV evolution phase. What differentiates this phase from the previous one is a greater focus on service agility, service latency performance, and implementation flexibility. To satisfy these demands, NFV implementations will need to embrace two additional attributes: distributed multi-cloud network support and workload portability achieved through software and hardware disaggregation.

Distributed Multi-Cloud Networks

Although the original NFV joint white paper introduced the concept of distributed NFV networks, the immediate need to scale the core meant there was a greater focus than on the edge. However, this is now changing. The first wave of edge initiatives, such as European Telecommunications Standards Institute (ETSI), defined multi-access edge computing (MEC) and 5G. Both adopt a fully distributed architecture to meet strict low latency performance budgets and, for some services, minimize backhaul bandwidth requirements.

While it is difficult to predict which services will dominate and drive service deployments even 6 years from now, there is little doubt that these services will be edge-based and leverage cloud-native intelligence in multi-cloud configurations.

Software and Hardware Disaggregation

Advanced cloud networks will also heavily rely on the principle of software and hardware disaggregation to complete the cloud transformation process. Software and hardware disaggregation, as shown on the right side of **Figure 1** below, involves the complete separation of key components, such as the operating system (OS), service logic, management tools, and even interworking protocols from underlying server hardware utilizing a software abstraction layer. The value of this approach is that it allows any service/workload to scale up and down on any hardware independently anywhere in the cloud.

In contrast, as illustrated in the middle part of the figure, in the initial NFV phase, virtualized software could run on commercial off-the-shelf (COTS) servers. However, retention of a monolithic software design meant that like the traditional integrated model shown on the left, it continued to be tied to vendor-specific entrenched OSs and management tools. This made hardware sharing and truly vendor-agnostic implementations problematic.

Figure 1: Integrated vs. Virtualized vs. Disaggregated Software/Hardware Design

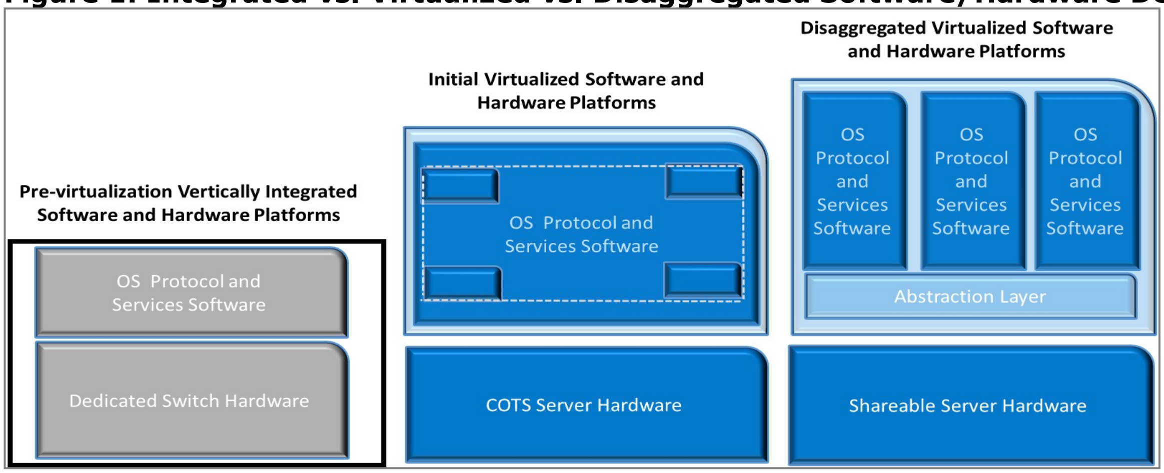

Source: Heavy Reading

SOFTWARE DISAGGREGATION

While some CoSPs may be willing to adopt an initial virtualized model because it simplifies implementation and accountability structure and provides some level of resource cost control, Heavy Reading believes that CoSPs will adopt a disaggregated model in large part due to the untapped power of software disaggregation. On a more granular level, the value of software disaggregation can be traced back to the support of several design characteristics, including the ability to support microservice software design using a common orchestration model in an automated and DevOps-centric environment.

Microservice Software Design

Disaggregation plays a major role in software service design. In this model, the approach is to break existing applications into smaller, modular microservices that can be reused and assembled independently for other new services. This is important because it aligns with the shared multi-tenancy application model inherent with cloud architectures.

This level of service decomposition not only provides greater flexibility in terms of creating and onboarding new services compared to monolithic software implementations, it also reduces service overhead since each microservice instance can be run in a software-based container instantiated according to current service load requirements. Minimizing overhead is important, especially with edge services, where resource capacity is limited. Therefore, 5G services have adopted a microservice design that can run agnostically on server platforms based on specific service requirements. The result is that for 5G and other edge low latency services, it is much easier to estimate end-to-end latency because the performance metrics of each microservice are constant and well understood.

The Power of DevOps

The adoption of software disaggregation also complements the shift that many CoSPs are undertaking to adopt DevOps software design principles. DevOps is attractive on several levels because it fosters the ability to use leaner software development and operations teams, and therefore improve service agility.

Microservices are highly reusable, so CoSP DevOps teams can use them to assemble new services independently without having to perform extensive code reviews with all software developer stakeholders. The latter is common practice with monolithic vertically integrated hardware and software platforms.

Implementing DevOps also simplifies and reduces costs associated with service development when using third-party software developers. This promises to dramatically shorten the time for new service development and deployment, which is traditionally often measured in years for CoSP services compared to the weeks or months of cloud development cycles.

Common Orchestration Model

Another drawback of the virtually integrated design model is that, traditionally, each vendor solution dictates which orchestration model(s) can be supported by the platform. Consequently, although NFV was designed to support a common orchestration model, implementation experience with initial virtualized systems has shown that it is difficult to implement. This not only introduces integration cost and implementation complexity constraints, but it also makes it challenging to implement open-source-based or even multi-vendor solutions.

With software disaggregation, since the software is reusable and most likely designed and implemented using DevOps principles, it is possible to fully achieve decoupling of software to hardware platforms. This translates into a more streamlined path to implementing a common orchestration model.

Automation Adoption

A final consideration is automation. There is no doubt that automation will ultimately become a staple of future networks because it can be successfully applied to a broad number of network layer functions.

Some of the key roles that automation will fulfill include support of real-time performance monitoring and security policy enforcement. However, given that each CoSP will also have unique automation requirements (depending on service and security strategies), maximizing automation implementation flexibility will be crucial. Software disaggregation will play a key role in enabling CoSPs to address these unique automation requirements using a broader range of ecosystem partners than is possible with initial NFV deployments.

The result of the combination of these advanced capabilities, as shown in **Figure 2** below, is the creation of a fully modular, disaggregated services model. This model not only fosters microservice reuse adoption, a common orchestration model, and seamless onboarding, it also supports cross-domain services that can run in a multi-cloud environment. Additionally, it puts in place a common reference architecture, which simplifies the introduction of powerful emerging capabilities, such as automation.

Figure 2: Implementing Modular Disaggregated Services

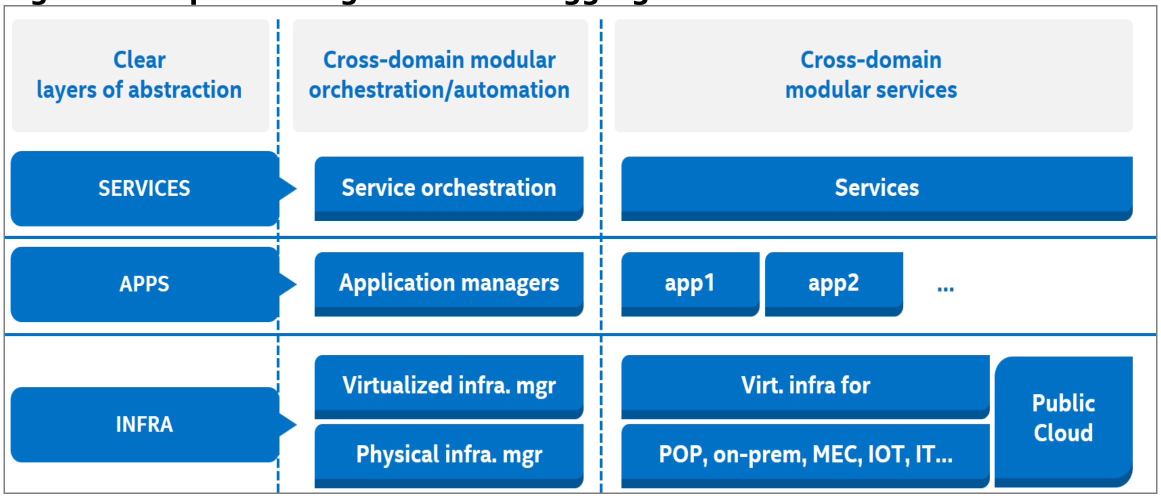

Source: Intel

Hardware Server Disaggregation

To fully realize the benefits of software disaggregation, hardware server disaggregation should also be supported. Hardware server disaggregation breaks down physical resource components into racks of pooled service resources that provide additional implementation flexibility. In this case, flexibility is derived from the ability to optimally share compute, input/output (I/O), and storage functions among other servers deployed in a cluster at the edge. In contrast, the current server model does not allow resource sharing, resulting in some server resources running at fully capacity while others are underutilized.

As summarized in **Figure 3**, the differences in terms of key attribute support between the initial virtualized platforms and disaggregated virtualized platforms are striking.

Figure 3: Comparing the Approaches

Attribute	Initial Virtualized Software and Hardware Platforms, 2012-2018	Disaggregated Virtualized Software and Hardware Platforms, 2019 and Beyond
Microservice support	Limited support – not typically supported	Fully supported – services are designed based on a microservices model
Distributed architecture	Limited support – not a primary focus	Fully supported – a primary focus area
Software disaggregation support	Limited support – reliant on vendor product design	Fully supported – a fundamental design consideration
Open-source support	Limited support – narrow ecosystem and vendor commitment	Fully supported – much broader ecosystem reach and vendor commitment

Attribute	Initial Virtualized Software and Hardware Platforms, 2012-2018	Disaggregated Virtualized Software and Hardware Platforms, 2019 and Beyond
Disaggregated platform support	Not typically supported – reliant on vendor product design	Fully supported – greater range of support from platform vendors
DevOps support	Limited support – due to vendor integration model	Fully supported – high degree of service reuse via microservices and open-source adoption
Automation support	Limited implementation flexibility – due to vendor linkages	Greatest level of implementation flexibility – due to service independence
Common orchestration model	Limited support – largely dependent on vendor capabilities	Fully supported – opening up the services layer enables adoption of a common orchestration model

Source: Heavy Reading

APPLYING SOFTWARE AND HARDWARE DISAGGREGATION TO 5G NETWORKS

This section of the white paper documents the benefits of applying disaggregation to both the distributed 5G new radio (NR) and next-gen core (NGC) network.

5G Distributed New Radio

One of the major benefits of the 5G NR is that it defines a fully distributed radio access network (RAN). This is accomplished by breaking traditional RAN radio components into radio units (RUs), distributed units (DUs), and a centralized unit (CU). In contrast while 4G remote radio units (RRUs) and the centralized baseband unit (BBU) support some level of separation they typically do not realize the vision of a fully distributed RAN.

Another way to look at this, as captured in **Figure 4** below, is that the 4G RAN supports virtualization aligned with the initial virtualized model, while the 5G RAN design methodology assumes hardware and software disaggregation in both the DU and CU.

Figure 4: Distributed 5G NR vs. 4G RAN

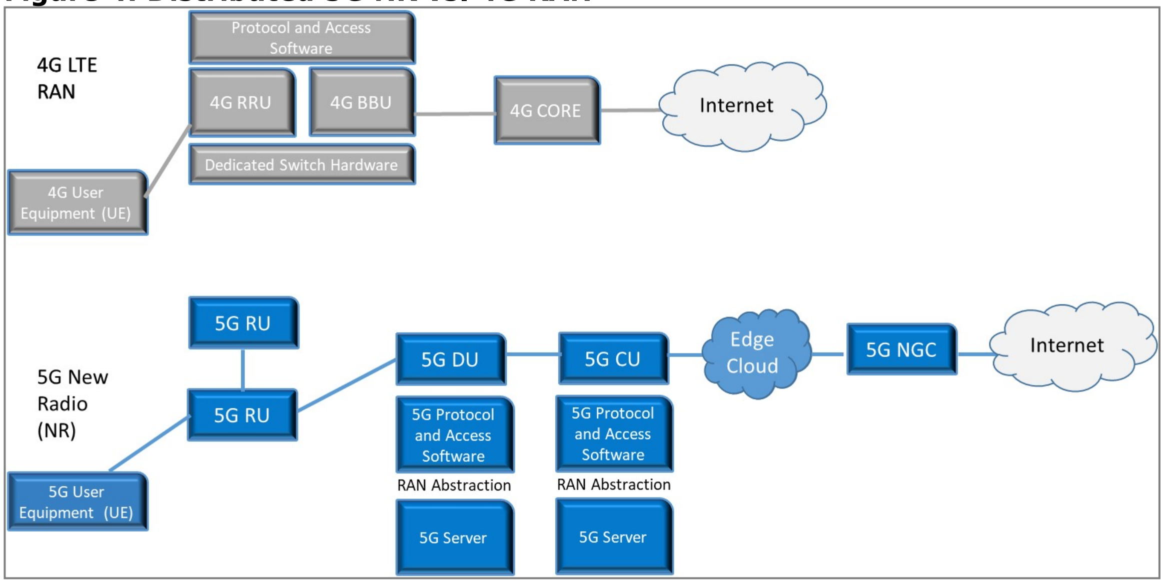

Source: Heavy Reading

This approach not only delivers the benefits of a disaggregated software model, it also maximizes resource scalability to most effectively manage radio VNF requirements. Additionally, it further unlocks the radio ecosystem, providing CoSPs greater vendor flexibility for independent selection of RU, DU, and CU components compared to the current vendor-centric integrated 4G RAN environment.

5G Distributed Next-Gen Core

Although CoSPs must deploy the 5G NR to launch 5G, implementing the corresponding new 5G NGC is optional.

Instead, CoSPs can choose to pair the 5G NR with their existing evolved packet core (EPC) and IP multimedia subsystem (IMS) core networks for control, policy, billing, and security support. This approach, referred to as the non-standalone (NSA) option, will be implemented by many CoSPs for initial 5G launch because it minimizes implementation complexity and minimizes new capital expenditure (capex) investment, while at the same time protecting existing network investment.

The trade-off with NSA is that it does not support the same reach of 5G-based sliced services or meet the performance levels necessary to support low latency, fully distributed microservices. Moreover, it means a continued reliance on the initial virtualized core model, which does not deliver the implementation flexibility benefits inherent with hardware and software disaggregation supported by the NGC. The benefits of applying disaggregation to both the centralized core cloud and edge cloud, as shown in **Figure 5** below, should not be understated.

Figure 5: Distributed 5G NGC vs. 4G Core

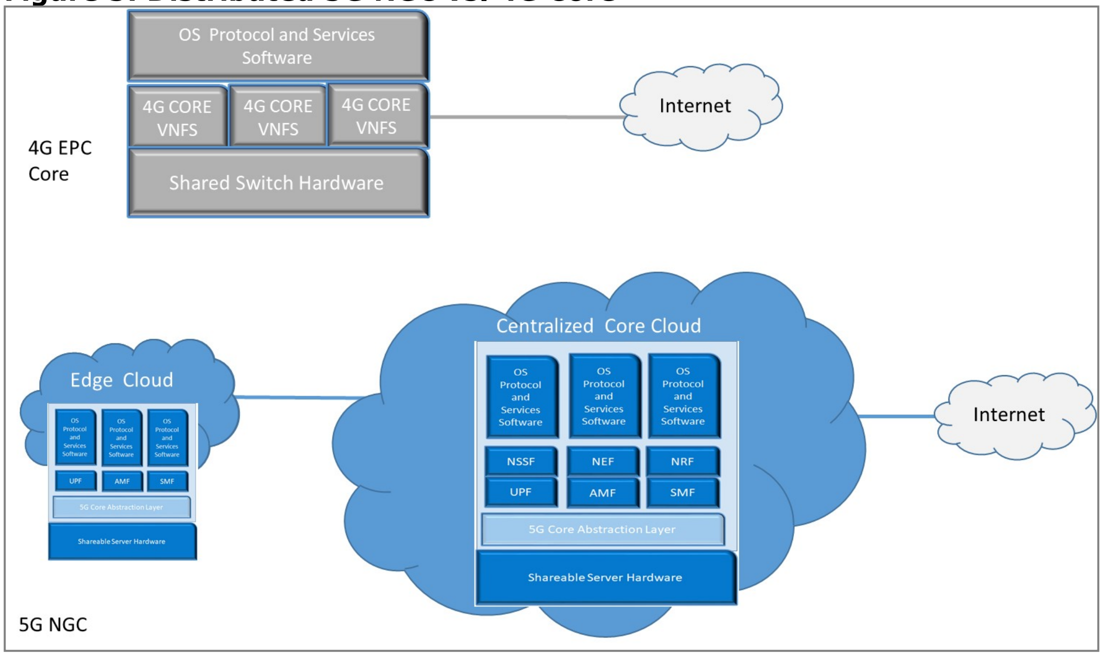

Source: Heavy Reading

Benefits like the 5G NR value proposition revolve around the ability to seamlessly deploy, provision, scale, and secure services anywhere in the network, using a pool of decentralized server resources. The end result is that CoSPs can leverage disaggregation to simplify delivery of telco services or enterprise services for adjacent market verticals, which is more complex with a 4G core.

In addition, implementing a disaggregated core also constitutes an important step in facilitating the adoption of heterogeneous orchestration models in an automated environment.

CONCLUSION

Over the past 6 years, virtualization has indelibly altered the telco landscape from both a technical and business perspective. As documented in this white paper, change is ongoing and driven by new demands associated with delivery of low latency microservices, distributed network architecture adoption, and the need to realize a common orchestration model.

Accordingly, CoSPs are now defining new requirements for their virtualized networks, which in turn is driving a new evolution phase. What sets this current evolutionary turn apart from the previous one is that it heavily relies on the principles of software and hardware disaggregation.

While even this current phase will likely still face implementation challenges, it is clear that disaggregating the virtualized network represents an important step in addressing these challenges. This is because disaggregation facilitates the creation and implementation of the flexible virtualization strategies necessary to holistically complete the cloud transformation.