HEAVY READING

WHITE PAPER

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

Network Slicing and 5G Future Shock

A Heavy Reading white paper produced for Ciena

ciena | blueplanet

JENNIFER P. CLARK, PRINCIPAL ANALYST, HEAVY READING

INTRODUCTION

The Future always comes too fast and in the wrong order. – Future Shock

2020 will mark the 50th anniversary of the publication of Alvin Toffler's futurist forecast of the impact of the digital revolution. *Future Shock* was published in 1970, when the digital revolution was in its infancy. Toffler was writing at a time of mainframes and punch cards – no disk storage and certainly no PCs or cell phones. And the smartphone was more than a generation away. Yet, his insights resonate as strongly today as they did 50 years ago, maybe more so.

In the intervening years, the industry has implemented four generations of mobile technology and is starting on the fifth. 5G promises 10 Gbps of throughput – 10 times the speed of 4G – and 10 ms of latency – 1/10th the latency of 4G. In addition to the new radio access network (RAN) technology – part of any new mobile generation – the mobile core will evolve to a new Service-Based Architecture (SBA). The changes do not stop here. 5G is now generally considered to include edge computing, cloud-native, container-based applications, and network slicing. While all these ancillary technologies can be implemented in 4G, the growing consensus is that they are mandatory with 5G. Together, these technologies will revolutionize the network at a pace that makes us question how the communications service providers (CSPs) will keep up.

In this white paper, Heavy Reading looks specifically at network slicing: what it is, why it is needed, what some of its dependencies are, and how the CSPs are likely to monetize it. What is clear from our research is that in order to succeed, 5G needs network slicing. In order for network slicing to succeed, it needs full life cycle orchestration and automation.

You've got to think about big things while you're doing small things, so that all the small things go in the right direction. – Future Shock

THE BEST THING SINCE SLICED BREAD?

A network slice is a logical end-to-end (E2E) network defined over virtualized resources on top of a common physical infrastructure (see **Figure 1**). Network slicing supports multiple unique virtual networks on the same physical infrastructure. It allocates specific resources to an application (e.g., an Internet of Things [IoT] app), a service (e.g., software-defined wide-area networks [SD-WANs]), a set of users (e.g., an enterprise), or a network (e.g., a mobile virtual network operator [MVNO]). These resources can be dedicated to one network slice only or shared between many. The network slices can be deployed across any part of a service provider network (e.g., edge and core) and can span multiple carrier networks.

Figure 1: Network Slicing

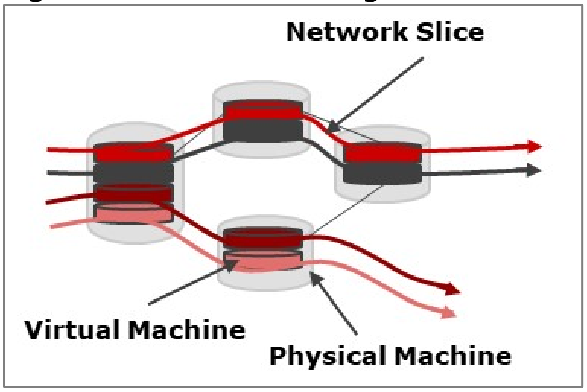

Source: Heavy Reading

Each network slice is designed according to the specific needs of the application or user, including speed, capacity, latency, security, and topology. Network slicing can be implemented on a 4G network. However, with the high speed and low latency promised by 5G, network slicing becomes essential.

With 5G, the industry will finally achieve one common mobile infrastructure that can support all categories of mobile traffic: voice, data, streaming video, cached video, gaming, IoT, or any combination thereof. However, how do we support different user requirements and traffic profiles without grossly overengineering the network, wasting infrastructure resources, or leaving customers vulnerable to security breaches? Virtualization, through the use of network functions virtualization (NFV) and software-defined networking (SDN), gets users part of the way there, improving network flexibility and how applications and endpoints are added, modified, or deleted. Network slicing on top of this virtualized infrastructure makes better use of network resources while improving customer experience.

The GSMA defines two sets of characteristics of a network slice: network connection services and network resource services. Network connection services refer to how the connection itself is defined. Network resource services are offerings apart from the network connection made available by the service provider to customize the offering to the enterprise. **Figure 2** shows a representative, but not exhaustive, list of both types of services.

Figure 2: Network Slicing – Infinite Diversity in Infinite Combinations

Network Connection Services	Network Resource Services
Low latency	Big Data analytics in support of management and orchestration
Stable and reliable high speed	ID and asset management, providing real-time automated authentication
Guaranteed SLA with service-level reporting	Platform security as a service
Guaranteed user experience across networks and/or nations	Storage as a service
Connected device management	Dynamic charging

Network Connection Services	Network Resource Services
Seamless mobility regardless of speed (e.g., high speed train), licensed or unlicensed technology (5G vs. Wi-Fi), or carrier	Cloud computing, providing access to the carrier's storage and compute resources
Energy efficiency (e.g., narrowband IoT [NB-IoT])	Edge computing, particularly for latency sensitive apps
Enhanced security	Partner integration for rapid integration of partner ecosystem virtual network functions (VNFs)

Sources: Heavy Reading, GSMA

Network Slicing Management and Automation

Network slicing will open many new revenue streams for the CSPs. Looking at the slice attributes in **Figure 1**, it is easy to anticipate CSPs designing and maintaining thousands of network slices. Vertical industry organizations such as the 5G Automotive Association (5GAA) have spun up work groups focused specifically on how to leverage network slicing. And they are not alone. Industry organizations such as Zentralverband Elektrotechnik und Elektronikindustrie (ZVEI) and Industrial Internet Consortium (IIC) are also investigating how 5G and network slicing solutions can contribute to smart manufacturing.

CSPs cannot provision and manage thousands of network slices manually or in a partially automated environment. As the implementation of network slicing accelerates, provisioning, life cycle operations, and configuration management must be dynamically orchestrated to achieve complete E2E automated services. Network slice orchestration and life cycle management have a long list of attributes they must assign, monitor, and manage, not just at setup, but continuously and in real time as network resources, traffic load, and service requirements change. The orchestration system must negotiate resources, connectivity, speed, latency, and security. These resources reside across different administrative domains in the network – RAN, transport, and core – and therefore point to the need for a higher level of abstraction. Resources may be dedicated, but most will be shared to create services that are cost-effective, yet high performance. At the same time, isolation between slices must be maintained. If the slice cuts across multiple service provider networks, information must be shared. Those resources must be orchestrated while protecting carrier autonomy and shielding sensitive or competitive information.

Determining how to design the management of network slicing cuts across multiple standards organizations. The 3GPP is the centralized entity working on the architecture, definitions, and functionality that will make up the Network Slicing Management Functions (NSMF) standard or set of standards. Other standards bodies contributing to NSMF include the following:

- **3GPP RAN:** Focusing on how RAN handles network slices.
- **ETSI Industry Specification Group for NFV (ETSI NFV ISG):** Focusing on making sure the network slice can run in the virtualized environment.

NSMF will also require a slice-specific view on fault, configuration, accounting, performance, and security (FCAPS) data and information from existing element and network management systems (EMS/NMS) on shared and/or dedicated network resources.

THE IMPACT OF 5G

When the conversation about 5G first heated up early last decade, there was widespread agreement about the goals of 5G: faster, denser, lower latency communications. However, unlike its 2G, 3G, and 4G predecessors, the new generation was not focused on a specific radio technology or the spectrum band(s), but rather, on creating a new user experience. By the end of 2018, the 3GPP declared that any network using 5G New Radio (5G NR) was, in fact, 5G. The NR technology, together with a set of spectrum bands, the use of millimeter wave technology, massive multiple input, multiple output (MIMO), and new beam forming technology, forms the foundation of 5G's technical portfolio. Lost in the sands of time, however, is that 5G was defined around the user experience – whether that demanded one or more of the following:

- Enhanced Mobile Broadband (eMBB)
- Ultra-Reliable Low Latency Communications (URLLC)
- Massive Machine-Type Communications (mMTC)

Network slicing has been key to 5G since its inception, bridging the 5G technology portfolio with its underlying principles for user experience. There are currently 23 5G carrier deployments globally. The majority are field trials or limited deployments, with South Korea being the standout. (The country's three mobile carriers claim over 1 million subscribers as of the middle of 2019.) In addition, there are 12 5G fixed wireless access (FWA) deployments. By 2023, Heavy Reading's sister company, Ovum, forecasts that 1.31 billion people will subscribe to the next-generation 5G network.

A network slice must be able to cut across all segments of the network – access, edge, and core; across multiple provider networks; and across management systems for these networks. As a result, over a dozen industry groups have a vested interest in creating and/or implementing network slicing standards (see **Figure 3**). Following is a brief description of some of the key responsibilities of the standards organization to give a flavor of how interrelated these standards are.

Figure 3: Network Slicing Standards Are Driven by a Large and Growing Ecosystem of Industry Groups

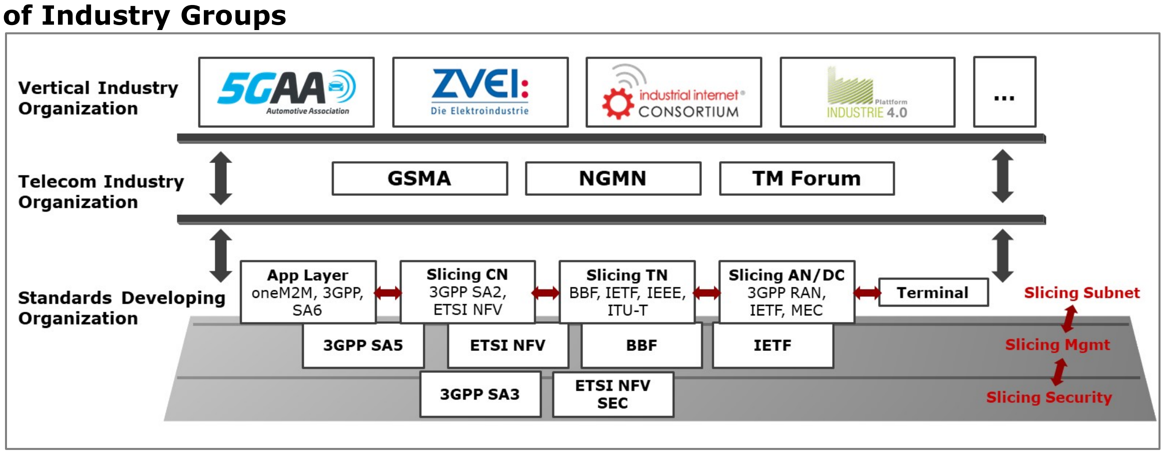

Sources: Heavy Reading, GSMA

3GPP

The 3GPP is at the forefront of standards development for network slicing and has multiple working groups involved in standards creation for network slicing, including the following:

- **3GPP SA1:** Focusing on network slicing use cases and their requirements.
- **3GPP RAN1/2/3:** Addressing network slicing in the RAN.
- **3GPP SA2:** Specifying the fundamental system architecture to support network slicing.
- **3GPP SA5:** Looking at E2E network slice creation and management within the 3GPP sphere of influence and also coordinating with other affected standards development organizations to generate complete E2E network slices.
- **SA3:** Addressing E2E network slicing security and working with ETSI ISG NFV on network slice isolation.

Broadband Forum

The Broadband Forum's (BBF's) role in network slicing is to define the transport network slicing architecture. It also works with the 3GPP to coordinate the interface requirements between the transport network and the 3GPP slicing management system.

IETF

Existing routing protocols such as segment routing and Layer 3 virtual private networks (VPNs) will be extended to support network slicing. The Internet Engineering Task Force (IETF) is also responsible for defining the interface between the transport network and the 3GPP slicing management system – working with both the BBF and the 3GPP to do so.

ETSI

There are two major groups within the European Telecommunications Standards Institute (ETSI) with network slicing responsibility:

- **ETSI ISG NFV:** Focused on resolving technical challenges of network slicing, such as computing and storage. This work group also determines how NFV technology supports network slicing, including standardizing the NFV slicing resource layer.
- **ETSI ISG ZSM:** The Zero Touch Network and Service Management Industry Specification Group (ZSM ISG) was created within ETSI in December 2017. The goal of the group is to resolve what is perhaps the most complex issue facing the technology: 5G E2E network slicing management. It aims to create a life cycle framework enabling fully automated delivery, deployment, configuration, assurance, and optimization of network services.

Industry Groups

In addition to the telecom standards organizations, there are also vertical industry groups that are drafting standards specific to network slicing, most notably the 5GAA, created in September 2016. By November of the following year, one of the work groups with the 5GAA had established a workstream to evaluate the potential impact of network slicing on the connected car.

Network slicing has tentacles into all the major network standards organizations. This should drive home to the carrier, the vendor, and the enterprise how critical it is to develop and implement standards-compliant network slicing solutions. Adherence to emerging standards is not only critical to seamless E2E connectivity for network slicing, but will also assist the CSPs in establishing a multi-vendor environment. Too often today, interoperability means going with a vertically integrated solution from a single vendor, locking the carrier into that vendor, and placing restrictions on how the network can scale horizontally in contrast with, for example, a cloud-native solution.

Cloud Native

Seven years after the concept of NFV was introduced, carriers are still not satisfied with today's VNFs in terms of functionality, ease of implementation, or management. Equipment providers have translated their physical appliances into virtual appliances, but few have taken the next step to cloud-ready applications and fewer still to cloud-native applications.

Cloud-native applications are packaged as lightweight containers and designed as loosely coupled microservices. From a carrier perspective, some of the most attractive aspects of cloud-native applications are their lower cost to develop, improved ease of upgrade and modification, ability to scale horizontally, and perhaps most importantly, the fact that they avoid vendor lock-in. While avoiding vendor lock-in is a fundamental goal of the CSPs, most of their vendor suppliers will move to cloud-native as slowly as competitive pressures allow. CSPs that are looking to move to cloud-native sooner rather than later will need to work with startups that have adopted a cloud-native approach from the beginning or with existing vendors that have moved to an open design philosophy encouraging multi-vendor deployments.

NETWORK SLICING USE CASES

Network slicing has plenty to offer the user in terms of network performance, security, isolation from others sharing the same physical infrastructure, and more. However, carriers will not implement the technology if they cannot monetize it. The carriers and their vendors should focus on coupling the deployment of 5G and network slicing, avoiding the stutter in adoption created when linked technologies are decoupled in terms of rollout and adoption – think Long-Term Evolution (LTE) and Voice over LTE (VoLTE). Network slicing will come at the price of software licenses and, in some cases, hardware modifications, plus the operational costs of deploying and maintaining it. The following is a set of use cases that Heavy Reading believes will benefit most from network slicing while creating a positive revenue stream for the carriers near term – over the next 2 to 3 years.

FWA – CSPs Are Anticipating the Opportunity

Carrier interest in FWA for 4G was high and has only increased with 5G, as evidenced by pre-standard implementations from carriers such as Verizon. Due to its use of millimeter wave spectrum and beam forming technology, 5G FWA can achieve download speeds of 25 Gbps or more – 10 to 20 times the speed of 4G LTE residential broadband. 5G FWA promises both speed and a low cost of deployment (compared to fiber to the home [FTTH]), making it the first technology that mobile network operators (MNOs) might use to compete effectively with the cable operators in the residential market. Network slicing would enable the carriers to provide differentiated services to residential and enterprise customers, including a low latency slice for gaming applications or a security slice for enterprises looking to protect their resources.

MVNO – Faster Route to Revenue

There are over 1,100 MVNOs globally, a number that will increase by an estimated compound annual growth rate (CAGR) of 4% between 2019 and 2024. This growth suggests that the business model for MVNOs is a compelling one. Unfortunately, this is not the case. The failure rate of MVNOs is staggeringly high, with 80% never making it a year beyond product introduction. The earnings before interest, taxes, depreciation, and amortization (EBITDA) for MVNOs averages in the red, unlike the MNOs, whose EBITDA hovers around 35% despite slowly declining revenue. Some of the market focus of today's MVNOs includes mobile broadband, digital content, digital commerce/payments, and IoT.

Whatever the application, the MVNOs are united in the need for efficient, flexible, and cost-effective infrastructure. Their infrastructure strategy starts with the need to pay their MNO partners only for the bandwidth used while at the same time having access to additional bandwidth when needed. A cloud strategy that leverages SDN, NFV, and network slicing is ideal for these purposes, enabling virtual instances to be spun up or torn down according to user demand. Nevertheless, sharing of the physical network resources is still complex and difficult to manage. Network slicing results in better isolation for the MVNOs. It also places control of the network slice in the hands of the MVNO itself, allowing it to define the service according to its own requirements – and not those of the host MNO's network. This results in a lower total cost of ownership, which enables MVNOs to launch faster and cut their time-to-revenue from scores of months to weeks and even days.

The Gaming Slice – Pokémon GO and Friends

The global video games market is forecast to expand at a CAGR of 5.6% over the next 5 years. By 2024, total revenue is projected to reach $197 billion. In recent years, growth has been driven largely by non-traditional gaming. Mobile games revenue from both end-user spending and advertising has expanded rapidly, especially in emerging markets.

Both residential and mobile gaming have the potential to be a driving use case for network slicing, particularly in combination with virtual reality and augmented reality. The high speeds and ultra-low latency of 5G FWA in combination with network slicing will result in a gamers paradise. And the gaming market, like the live streaming of sporting events, shows reduced price sensitivity.

Augmented reality headsets and glasses have not seen widespread popularity due to the tug of war between user comfort/ergonomics and the processing power needed on the headset to pull off cool effects. With 5G and network slicing, headsets could be similar to a pair of sunglasses, connected via the network slice to an edge computing platform. The platform could perform any needed computations either locally or, if needed, in the cloud and then, at high speed and low latency, return results to the headset.

Smart City – This Market Is Gaining Momentum

Heavy Reading defines a portfolio of seven key smart city applications:

- Smart lighting
- Smart buildings
- Smart urban transport
- Smart parking
- Public security
- Traffic management
- Environment management

Low power wide-area networks (LPWANs) are enabling new smart city use cases. LPWA can be 5-10 times cheaper than traditional cellular; thus, vendors and CSPs can be more creative in apps that they bring under the smart city umbrella. Smart city rollouts are accelerating, particularly in the Middle East and East Asia. Digital tags – such as for parking – mean that the consumer shoulders the cost burden for the application, certainly initially.

Sales of smart city IoT devices will grow at an estimated 27% CAGR between 2018 and 2022 to nearly 600 million devices. Most of these sales will be in low power, low data consumption devices. The benefits of network slicing in this use case are readily apparent. Public safety and first responders will be allocated a slice that is defined as high priority and highly reliable. IoT applications can be allocated their own slices using, for example, NB-IoT. Edge computing can be utilized for applications that require improved response time and would benefit by keeping more traffic local. Privacy with some apps can be established with isolated network slices and enhanced security attributes.

Figure 4: Low Power Radio Modules for IoT Are Priced to Sell

Technology	Module Cost
Sigfox	$1-2 (single radio)
LoRaWAN	$5-6 Nano-Tag $2
NB-IoT	$5 retail price (T-Mobile USA)
LTE-M	$7-10 retail price (AT&T, Verizon)

Source: Informa Tech

Alongside the above use cases will be the ubiquitously discussed connected car apps and, eventually, autonomous driving. There is other low hanging fruit that the service providers will rotate into their portfolio. Consider, for example, the perennial security and management issues of bring your own device to work, whether laptops or mobile phones. With network slicing, a single user device can belong to two network slices: one with a personal profile and one that connects securely to the corporate network.

The 5G and network slicing markets are not emerging overnight. The standards, management, orchestration, and automation are still under development. CSPs with large software development capabilities are trialing these services today, but it will be 2 to 3 years before the carrier market begins to implement network slicing along with 5G.

CONCLUSIONS

> *The illiterate of the 21st century are not those who cannot read and write, but those who cannot learn, unlearn and relearn.* – Future Shock

The migration to 5G brings with it tributary technologies that together are transforming the network: cloud-native applications architected with containers and microservices, automated service life cycle, edge computing, and network slicing. These advances in network technology will eventually affect every aspect of how people and their possessions communicate. Successful implementations will be standards-compliant and will leverage a cloud-native architecture to ensure agility and scalability and avoid vendor lock-in. The network slicing solution will be fully automated and life cycle managed in order to deploy new slices on demand and manage and upgrade existing ones. CSPs should be trialing multiple solutions and working to avoid vertically integrated single vendor solutions. They should opt for a modular solution that allows the CSPs to go with best-in-breed components in a multi-vendor architecture.

ABOUT BLUE PLANET

This section was written by Blue Planet.

Blue Planet, a division of Ciena (NYSE: CIEN), provides market-leading intelligent automation software and specialized professional services to help clients modernize their IT and network operations. The Blue Planet Intelligent Automation portfolio provides a comprehensive and open software suite that lets service providers and mobile network operators (MNOs) use deep knowledge about the network to power the intelligent transformation and adaptive optimization of their services and operations.

Deployed by major network providers worldwide, Blue Planet's diverse portfolio of products and solutions are designed to help operators accelerate their transition to 5G while reducing onsite engineering requirements and associated costs. The Blue Planet software suite includes solutions to automate the orchestration and activation of E2E network slices, seamlessly scale the bandwidth requirements of your 5G deployments, optimize the network to enable better load spreading, and assure customer experiences are continually optimized for their varied requirements.

The combination of the Blue Planet Intelligent Automation software with Ciena's programmable packet-optical network solutions allows MNOs to realize "Zero Touch" planning, design, and operations capabilities to ensure the reliable delivery of new 5G services, with minimal manual intervention in the shortest possible time.

For updates on Blue Planet, visit www.blueplanet.com. To learn more about Ciena 5G Network Solutions, visit www.ciena.com/insights/5g.