Building the Network of the Future

Getting Smarter, Faster, and More Flexible with a Software Centric Approach

Edited by John Donovan and Krish Prabhu

CRC Press
Taylor & Francis Group
A CHAPMAN & HALL BOOK

Chapter 2 – Transforming a Modern Telecom Network – From All-IP to Network Cloud

Rich Bennett, Steven Nurenberg

This chapter provides context for and introduces topics covered in subsequent chapters. We summarize major factors that are driving the transformation of telecom networks, characteristics of a modern all-IP network, and briefly describe how the characteristics change as an all-IP network transforms to a network cloud. The brief descriptions of change include references to succeeding chapters that explore each topic in more detail.

2.1 Introduction

The dawn of the Internet began with industry standards groups such as the International Telecommunications Union (ITU) and government research projects such as ARPANET in the US proposing designs based on loosely coupled network layers. The layered design combined with the ARPANET open, pragmatic approach to enhancing or extending the design proved to have many benefits including: integration of distributed networks without central control; reuse of layers for multiple applications; incremental evolution or substitution of technology without disruption to other layers; and creation of a significant sized network in parallel with the development of the standards.

Tremendous growth and commercialization of the network was triggered by: the addition of a simple TCP/IP application Hypertext Transfer protocol HTTP; availability of a light weight application for commonly used UNIX servers; and availability of a navigation and information display application for the installed base of personal computers. This convergence of events created the World Wide Web where islands of information could be linked together by distributed authors and seamlessly accessed by anyone connected to the Internet. The Public Telecom network became the transport network for internet traffic. Analog telephony access circuits were increasingly used with modems for internet data, driving redesign or replacement of access to support broadband data. Traffic on core network links associated with internet and IP applications grew rapidly.

The IP application protocols were initially defined to support many application types such as email, instant messaging, file transfer, remote access, file sharing, audio and video streams, etc. and to consolidate multiple proprietary protocol implementations on common network layers. As the HTTP protocol became widely used across public and private networks usage grew beyond the original use to retrieve HTML Hypertext. HTTP use today includes transport of many IP application protocols such as a file transfer, the segments of a video stream, application to application programming interfaces, and tunneling or relaying lower level datagram packets between private networks. A similar evolution happened with mobile networks that started providing mobile voice and today with 4G LTE wireless networks, IP data is the standard transport and used for all applications including use for voice telephony. IP has become the universal network traffic standard and modern telecom networks are moving to all IP.

2.2 Rapid Transition to All-IP

The convergence of technologies to all-IP is driving major changes in telecom networks. The price and performance of mobile computing devices and the emergence of a vibrant open software ecosystem is creating an exponentially increasing demand for mobile applications and data access. Telecom service providers have made large investments to meet mobile applications demand with wireless network technologies such as 4G LTE having embraced the use of IP as a foundational architectural element. The 3GPP standards provided a blue print for integrating the existing telecom services and future IP service with an IP Multi-Media Subsystem (IMS) and alternatives for IP services platforms have emerged from adjacent industries. Public compute, storage, content delivery clouds accessed via IP protocols have changed the transport usage patterns for both business and consumer users. Business applications leverage shared resources in the cloud rather than private, distributed data centers and consumers increasingly interact more with the content in the public clouds. Today, a new telecom network, deployed in a "greenfield" fashion, will be exclusively designed, built and operated in an all-IP manner end-to-end.

2.3 The Network Cloud

The declining significance of transport and emergence of an all-IP network platform with diverse IP services demand creates an opportunity, or in some cases an imperative for telecommunications providers, to remain relevant by embracing a new framework for the delivery of services. The essential elements of the new framework include: refactoring hardware elements into software functions running on commodity cloud computing infrastructure; aligning access, core, and edge networks with the traffic patterns created by IP based services; integrating the network and cloud technologies on a software platform that enables rapid, highly automated, deployment and management of services, and software defined control so that both infrastructure and functions can be optimized across change in service demand and infrastructure availability; and increasing competencies in software integration and a DevOps operations model. We call this the **"Network Cloud."**

The benefits of a Network Cloud framework include lower unit cost of providing services, faster delivery and flexibility to support new services compared to the traditional approach (where a specific service requires dedicated infrastructure with low utilization to ensure high availability), and the opportunity to automate network operations.

2.4 The Modern IP Network

The modern network as described in more detail in Figure 2.1, is formed using high capacity fiber optic transmission, optical switching & multiplexing, and packet switching technology. Each of these is combined to create local, regional, national, and international networks. In the past, virtually all telecommunications companies tended to specialize in both service type and geographic reach. With newer technologies, it has

become possible to use the same network to offer a range of voice, video, data, and mobile services leading to significant competition and consolidation.

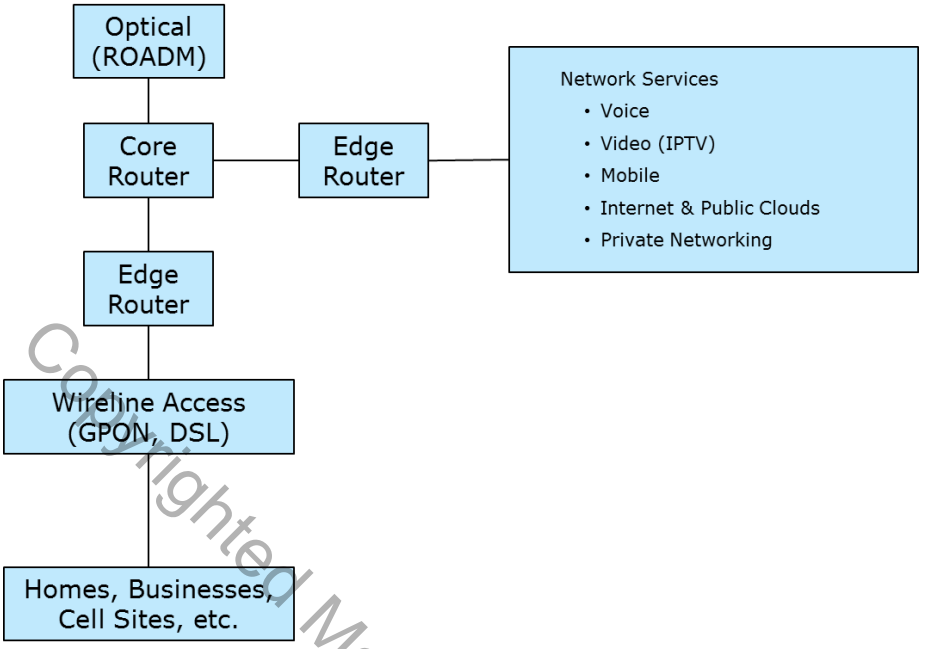

Figure 2.1 – Illustration of a Modern IP Network

2.4.1 The OSI Model

The primary model for networking used today is the Open Systems for Interconnection (OSI) model which breaks network functionality down into layers (see Figure 2.2). This allows technology enhancements and change to occur without disrupting the entire stack (combination of layers). It also allows inter-operation between networks (or network segments) that use dissimilar networking technology as long as they connect using a common interface at one of the layers. This is the method that allowed the Internet and its lower layer data networks to evolve so quickly. Each of the different local area networks, or LANs, could evolve as long as they all connected using the common Internet Protocol.

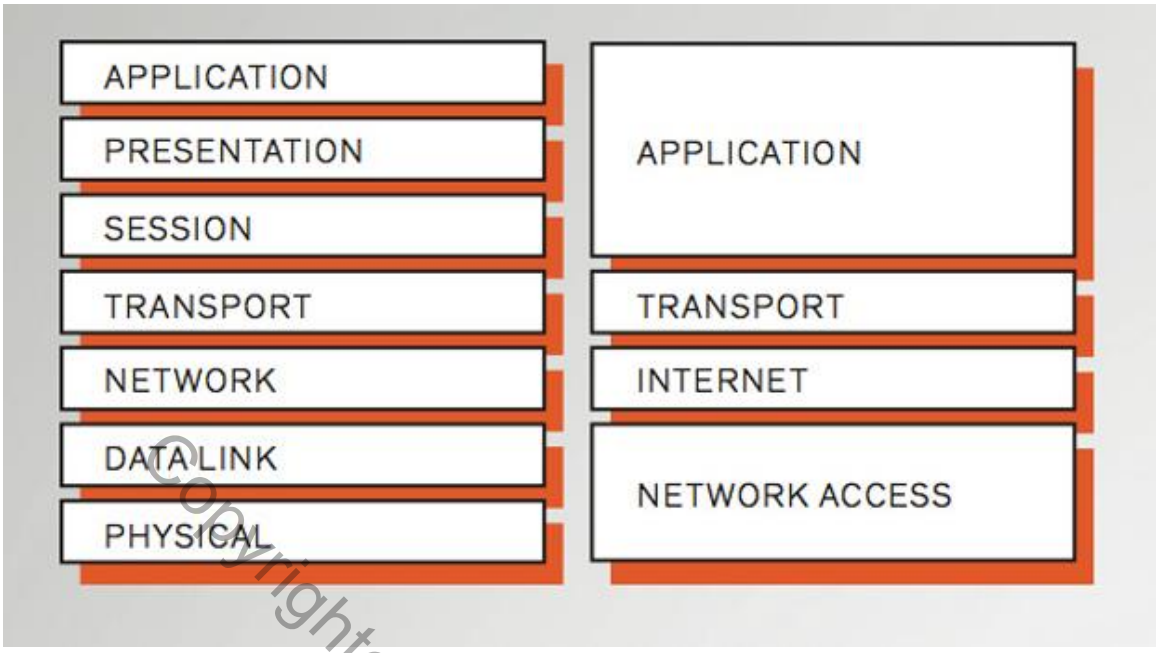

Figure 2.2 – OSI and TCP/IP Stacks and Layers[1]

2.4.2 Regulation and Standards

No matter how organized, telecommunications networks in order to provide value, need to be built against a set of standards. Internationally, the United Nations has historically had an arm called the International Telecommunications Union (ITU) that ensured the creation of standards based telecommunications networks. With the creation of newer technologies, additional technology groups or consortia were created to facilitate the advancement of specific approaches. For the Internet, this is the IETF (Internet Engineering Task Force); for mobile networks, this is the 3rd Generation Partnership Project (or 3GPP).

To help put the transformation of public networking into context, this section describes two examples – a traditional "appliance" based network and a contemporary software-defined and virtualized network (SDN/NVF) at a high level. This is to convey that that the new SDN/NFV network is more than just faster networking and certainly different than previous efforts to add intelligence and programmed control.

2.5 Transforming an All-IP Network to a Network Cloud

Since every part of a modern telecom network can be architected to be all-IP, it can also be transformed to a Network Cloud architecture. This transformation needs to be addressed individually for each of the components of the network – customer premise,

[1] spectrum.ieee.org/computing/networks/osi-the-internet-that-wasnt

access, edge, core, and service platforms. This is because, the benefits of NFV and the degree of SDN control varies from part to part. But there is a great deal of commonality – the NFV infrastructure, the SDN control framework, cloud orchestration, network management, security considerations, data collection and analysis, and a new operations paradigm for the software controlled network. This illustrated in Figure 2.3 where traditional routers are replaced by a network fabric and where the various dedicated elements are replaced by software running on COTS servers to implement the data, control, and management planes.

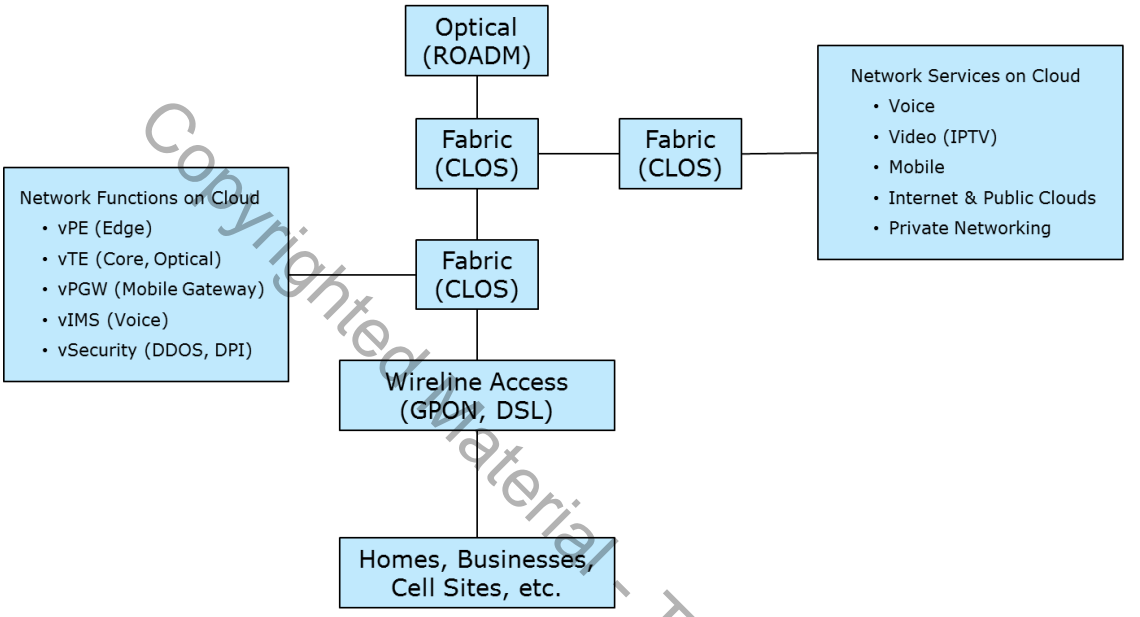

Figure 2.3 – Transformation to Network Cloud

2.5.1 Network Function Virtualization (NFV)

Network Functions Virtualization (NFV) draws equally from Information Technology (IT) and Network. Over the past decade, IT environments have gone through multiple stages of evolution, with the most recent being the use of virtualization that allows multiple software environments to share the same compute and storage hardware. The Network aspect comes about when this same technique is applied to functionality used in networking.

In the past, the networking functionality was created in purpose-built equipment that contained the hardware and software bundled together to meet the customer's specifications. Examples include switches/routers (OSI Layers 2 and 3), firewalls, session border controllers (used for voice and media networking), and mobility management entities (MME, used for signaling in mobile networks). Depending on the equipment, the hardware ranged from being a basic server, a server with special cards, or a custom designed element using proprietary ASICs (Application Specific Integration Circuits). In most cases, the underlying hardware can be separated from the operating software through an abstraction layer and the hardware platform deployed could be

common off the shelf hardware (COTS), which has very low unit cost because of widespread use across many industries. However, networking does have a requirement for special purpose hardware, either for specific physical interfaces such as optical or radio, or for throughput reasons (exceeding 10 Gb/s). In these, functional separation can be used to minimize the use of specialized hardware.

NFV leverages the IT environment, specifically the cloud paradigm, separating network software from hardware and leveraging the fact that today's modern servers are sufficiently scalable for most network functions where logic is employed. NFV related work accelerated when a group of network operators collaborated with each other on requirements establishing a study group under the auspices of the European Standards Institute ETSI (cite ref at end of chapter - http://www.etsi.org/technologies-clusters/technologies/nfv). This group was established in November 2012 and quickly gathered industry wide participation from carriers, vendors, chip designers, academic institutions, and other standards bodies. It laid out the logical framework for NFV and the critical functional blocks along with their interaction and integration points. Transforming physical network functions into virtualized network functions introduces a host of challenges which are discussed in more detail in Chapter 3.

2.5.2 NFV Infrastructure (NFVI)

Using the NFV approach, traditional network elements are decomposed into their constituent functional blocks. Traditionally, in a telecom network built with appliances, software is used in two major ways - network element providers acquire and/or develop the software components to support the operation and management of the network element, and service providers acquire and/or develop the software components needed to integrate the network elements into business and operational support systems. However, in a NFV model, the potential for greater efficiency is expected from more standardization and reuse of components that go into the network functions and as a result, the effort to integrate functions and create and operate services is lower because they can be supported in a common NFVI (NFV Infrastructure) platform.

Software plays a very important role in the transformation of a traditional telecom network to a Network Cloud. The software discipline provides extensible languages and massive collaboration processes that enable: solutions to complex problems; reusing solution components to automate tasks and find information and meaning in large amounts of raw data; and insulating solutions from current hardware design in a way that allows solutions to evolve independently of hardware and to benefit from a declining cost curve of commodity computing infrastructure.

In the new architecture, physical network functions (PNF) are supplanted by virtualized network functions (VNF) that run on NFVI, which is the heart of the Network Cloud. Key drivers for the platform are: the ability to utilize open source software, create sharable resources, and provide for the use of modern implementation practices like continuous integration / continuous development (CI/CD). The Open Stack environment is an obvious choice for NFVI - Open Stack provides the functionality to schedule

compute tasks on servers and administrate the environment. Additionally, it provides a fairly extensive set of features; this is discussed in more detail in Chapter 4.

2.5.3 Software Defined Networking (SDN)

While the general perception is that SDN for networks is a relatively new phenomenon with roots in data centers, the reality is that telecom networks have been using these techniques for several decades. An example is the use of techniques to control the data plane in circuit-switched networks, via the control plane from a controller, to provide 800 number services (for long distance automatic routing of calls and reverse billing).

In the Network Cloud, SDN is further advanced by applying these learnings and those from cloud data centers, to create a layered distributed control framework. A "master" (or global) controller which has a network-wide view of connectivity, controls subsidiary controllers associated with different functional parts of the network. Each controller may also be replicated for scale. For the IP/MPLS network fabric, the Network Cloud uses a hybrid approach retaining distributed control planes and protocols for responsiveness and using the centralized control plane for optimization and complex control.

One very unique aspect of Network Cloud SDN is the shift to model driven software and network configuration. In the past, these tasks were done by creating detailed requirements documents that network engineers then had to translate into the specific vendor configuration language of the associated network element. Any introduction of new network capabilities took extra time waiting for documentation, testing configuration, and implementing the configuration in operational support systems (OSS). The new process involves defining capabilities using the YANG (Yet Another Next Generation) modeling language[2]. YANG templates, which are vendor neutral, create a portable and more easily maintained definitive representation of desired network functionality. Another benefit is that simulation can be used to verify correctness before using the templates with actual network functionality.

In AT&T's approach to SDN, all SDN controllable elements are treated similarly, whether they are implemented as classic physical network functions (PNF) or the more modern virtual network functions (VNF). The Network Cloud SDN controller has a number of subsystems and is designed to operate in conjunction with ONAP platform. It combines a service logic interpreter (SLI) with a real-time controller based on Open Daylight (ODL) components. The SLI executes scripts that define actions taken on lifecycle events such as service requests and closed loop control events. SLI frequently used or complex scripts can be encapsulated as a java class that runs in the ODL framework then as reused as a single script operation. Another function is Network Resource Autonomous Control which is used to associate (assign) network resources with service instances.

[2] https://tools.ietf.org/html/rfc6020

At the bottom of the controller are adapters that can interact with a wide-variety of network element control interfaces These can range from older element management system (EMS) style “provisioning” to real-time network transactions using BGP.

No control framework would be complete without a policy function. Policy rules can be defined based on events collected from NFVI and services running on it. Events can trigger changes in algorithms at enforcement points in and/or control actions on the NFVI/SDN network cloud. For example, a high utilization of a network link might trigger a traffic engineering routing change or an unresponsive service function might trigger a restart of the service function.

Using the advanced SDN controller capabilities, it is possible to create not only conventional network services like Internet, Virtual Private Networks, Real-Time Media Services, etc., but also more complex on-demand services. Examples described in more detail in Chapter 6 are VPN service ordering, Bandwidth Calendaring, and Flow Redirection.

2.5.4 Open Network Automation Platform (ONAP)

The movement to network function virtualization changes many aspects of the way network infrastructure and services are managed over the life cycle of operation. The initial design of a service is no longer a vertically integrated optimized infrastructure, but must assume use of a distributed cloud hardware infrastructure and re-use network functions from any source that best meets the needs in the service design. The initial installation, configuration, turn up, and then changes in response to lifecycle events must be formally described in ways that can be automated such as defining work flows and actions that reference run-time sources of information and adjust use of infrastructure resources. The infrastructure as well as the service functions must expose software interfaces that allow monitoring, external control, and specification of policies that change run-time behavior. Both short term traffic engineering and long term capacity planning decisions must consider a broad range of services and scenarios making use of the common infrastructure.

Traditional Operations Support Systems (OSS) and Business Support Systems (BSS) were designed to integrate monolithic network element components and to add capabilities necessary to deliver and support the operation of customer services. This approach has a number of limitations. Network element components lack standard interfaces for some lifecycle operations and tend to be optimized for a particular service or not easily shared across different services. This increases cost and time to perform lifecycle operations such as initial delivery, upgrades, routine maintenance, repair of faults, etc. and requires dedicated infrastructure and skills. The time to design, integrate, deploy infrastructure with staff trained to operate it at scale limits the flexibility of a service provider to deliver new services. New or emerging service volumes and uncertainty of growth make it hard to justify the investment and the long lead time to deliver increases risk of missing market opportunities.

The ETSI NFV[3] effort described above produced a specification for NVF Management and Orchestration (MANO). ONAP expands on ETSI MANO by adding: a comprehensive service design studio to on-board resources, create services, and define life cycle operations; closed loop control based on telemetry, analytics, and policy driven life cycle management actions; a model driven platform to accelerate and reduce costs of on-boarding functions and eliminate VNF specific management systems; and support for physical network functions.

A consistent plan that includes both how ONAP capabilities replace traditional OSS/BSS systems such as fault correlation, performance analysis, and element management and how these traditional systems are phased out is critical for the transition period where there is a mix of D1 and D2 infrastructure. Without a clear plan for both there is a risk that traditional systems and design assumptions get integrated with service designs on the ONAP platform thus decreasing the benefit of highly automated operation and increasing maintenance costs.

Chapter 6 describes the ONAP software platform that supports the design and operation of services in a Network Cloud. This platform is visible as: APIs to software designers and developers creating services and virtual functions; operations to view data exposed on the run-time state, events, and performance; to Business and Operational Support Systems that interact with real-time events; and customers in using, higher level interfaces where they need to configure services and integrate with their private infrastructure.

2.5.5 Network Security

Security for networks is a multi-disciplinary problem – starting with the classic security problems such as confidentiality and integrity, coupled with the problem of availability, i.e., ensuring the service and network work, and can withstand a hostile attempt to bring it down. With Network Cloud, there is the need to understand security through the combined lens of software and network, in an operating cloud environment. For example, using the security technique of "defense in depth" and "separations of concerns," it is a pragmatic approach to categorize virtualized network functions keeping each category type from running on the same server simultaneously. For these and other reasons, security is a key architecture and design factor for the Network Cloud.

Security manifests itself from two different perspectives – the protection of the infrastructure itself and the ability to delivery security functionality in a service. The basic structure for infrastructure security is to ensure each component is individually assessed and designed with security in mind. For the Network Cloud, the fabric component contains functionality like access control lists (ACLs) and virtual private networks to limit and separate traffic flows. The servers run operating systems with hypervisors that provide separate execution environments with memory isolation. Around this, operational networks that are used to provide access to management ports, use firewalls to provide a point of control and inspection.

[3] http://www.etsi.org/technologies-clusters/technologies/nfv

Other aspects of security architecture also come into play. While the desire is to prevent a security problem, a good security approach also provides mechanisms for mitigation, recovery, and forensics. Mitigation approaches for infrastructure include overload controls and diversion capabilities. Overload controls (which also help in the case of network problems) prioritize functionality so that resources are best applied to control plane and high priority traffic. Diversion is the ability, typically using specific portions of the IP header to identify candidate network traffic, to re-direct packets for subsequent processing, rate limiting, and when necessary, elimination.

Forensics is fundamental to security. The ability to log activity provides the ability to analyze past incidents to determine root cause and to develop prevention and mitigation solutions. It also plays a role in active security enforcement by acting as an additional point of behavior inspection which may indicate either the failure or a weakness in the security design or implementation.

Within security processes, two key components are automation and identity management. Automation allows for the administration of complex sequences and eliminates the human element from configuration that is a typical source of security problems. All it takes is entering the wrong IP address in an ACL to create a vulnerability. Identity management ensures both people and software are authorized to view, create, delete, or change records or settings. The Network Cloud uses a centralized approach for identity verification. This prevents another weakness of local passwords which are more easily compromised. These topics are discussed in more detail in Chapter 9.

2.5.6 Enterprise Customer Premises Equipment (CPE)

Enterprise networks are the environment that businesses use to tie together their people, information technology, and operating infrastructures. From offices, warehouses, factories, and on the go people in trucks, cars, and when visit their customers, enterprises need comprehensive solutions for communications. Typically, they use a range of voice, video, data, and mobile services. All of their operating locations need some form of on-premises equipment, called Customer Premises Equipment (CPE.) In the past, CPE followed the same approach as network equipment, i.e., implemented as appliances. However, CPE is undergoing the same transformation using NFV. This allows a single device to provide multiple functions when needed, under software control.

When re-designing CPE using NFV, the approach taken was to allow functions to operate either on the premise inside the new virtualized CPE or in the network, on the Network Cloud. For the on-premises network functions within the CPE, innovation was required to create a suitable execution environment. Unlike the network cloud where multiple servers are available to scale up or down VNF, the CPE environment is limited to typically a single CPU chipset. To allow for software portability and to leverage the open source ecosystem, again KVM was selected as the hypervisor to allow multiple VNF to share the CPU in individual virtual machines. (There are less security concerns here since CPE is dedicated to a single customer.)

One of the most challenging aspects for CPE is management. Since CPE is located on-premises at the point between the end of a customer's wide area network (WAN) service and the local network, it is important to provide an operational network connection capability. This is done by sharing the WAN service segregating the traffic with special VLAN or IP addresses. However, before the service is initiated or during service after a failure, the WAN connection may not be available. The solution was to leverage mobile access to the network cloud and provide a "call home" capability. This second connection allows for remote "test & turn-up" procedures and in the case of WAN failure, diagnosis and fault location determination. This is discussed in more detail in Chapter 10.

2.5.7 Network Access

Access is used to describe the portion of the network that goes from homes and businesses to network infrastructure buildings called central offices or points of presence. This "last mile" of the network is also the most expensive investment a carrier must make since it provides connectivity across the geography of the service area involving the installation of media and equipment termed "outside plant". Historically most access was based on copper twisted-pair cabling that followed the roadways outwards. Modern networks use fiber as the preferred media for "fixed" or "wired" access with the latest approach called Gigabit Passive Optical Networking (or GPON). (The older metallic copper or for cable networks, coax, is still used into the home or business except where a rebuild has happened that delivers fiber all of the way as a "fiber to the home" or FTTH.) With the emergence of mobile (cellular) communications as a method for virtually anywhere communications, many have decided to forgo traditional wired service and go completely wireless. Like fixed networks, these networks have also undergone major changes, the most recent being the Long Term Evolution (LTE) or 4G (for 4th generation) which is based on an all IP structure.

The goal of the access network is to provide each customer with reliable and high performing connectivity to the network in as economical a manner possible. This is done by maximizing the amount of the infrastructure that can be shared. For GPON, sharing is done by combining up to 64 service connections onto the same single strand of fiber. For LTE, it is done by sophisticated radio control protocols that shares radio spectrum as efficiently as possible.

With the network cloud, access technology is being fundamentally transformed by leveraging the distributed network fabric, compute and storage capabilities to host access related network control and management functions. The data plane of access, due to its unique aspects, will continue to the use of specialized hardware that needs to be very close to the access media, but even here it is envisioned that programmability within that hardware can be brought to bear to provide for flexible evolution. Other access functions can run in the network cloud providing management and control functionality, but must typically be located nearby (e.g., under 10 miles) in order to be able to operate in the necessary timescales.

Chapter 11 gives a detailed overview of modern access technology and how it is being transformed using SDN and NFV technologies and the network cloud.

2.5.8 Network Edge

In packet networks, the edge platform is where network connectivity services are delivered. Before the advent of SDN and NFV, this was done using specialized network elements that converted and multiplexed the bit streams of customer access onto inter-office trunks while applying service functions. With IP networks, this is done using a router which implements the Provider Edge (or PE) functionality. A router was designed in a chassis with process cards to handle control and management functions, fabric cards that interconnected cards across the chassis, and line cards that could be used to connect to customers or to other portions of the network (typically PE connects to P core routers). The PE is configured for the type of service purchased such as Internet or Private Network.

With the network cloud, two transformations occur. First, physical connectivity is shifted onto a fabric built using merchant silicon that offers higher density and lower cost than the proprietary and specialized PE routers. Second, routing functions are disaggregated with control and management shifted to software on the network cloud. Third, the customer data plane is split between the fabric and software switching that runs as software. This last part allows a variety of options for providing PE services. For example, instead of deploying a proprietary PE dedicated to Internet or Private Networking, multiple PE software elements can run within the same server and provide the same services. If logical resource limits or specialized configurations are required, additional PE software can be executed and dedicated.

The network edge is described in more detail in Chapter 12.

2.5.9 Network Core

Chapter 13 describes in more detail, core networking for optical (Layer 1) and packet (Layers 2 & 3) in the OSI layer model. The optical layer provides fiber optic interconnect between remote central offices; this entails electrical to optical conversion, optical amplification, transport controls, and multiplexing/de-multiplexing of multiple signals. The packet core acts as a “backbone” that ties together the periphery allowing packet traffic to enter the network virtually anywhere and be economically forwarded to its intended destination.

The electro-optical equipment used for modern fiber communications falls into two basic types: the reconfigurable optical add/drop multiplexor (ROADM) and the optical amplifier (or OA). The former is placed typically in operator central offices to allow traffic to be added and removed (similar to on and off ramps of the highways). The latter is responsible for increasing the strength of the light so that it can continue on towards its destination.

Modern high capacity optical systems work by combining a number of separate high-rate (currently 100 or 200G, but soon to be 400G) optical signals onto the same fiber, using a technique known as Dense Wavelength Division Multiplexing (DWDM) where each signal is assigned a unique wavelength. The ROADM adds onto this by providing the ability to add or drop specific wavelengths at each node. This forms an optical network allowing a large number of locations to send optical wavelengths between each other without requiring separate fiber optic cables

With the advent of software defined-networking, the optical layer is being further advanced in two key ways – disaggregation of functionality and global optimization of transport resources in conjunction with the packet layer. Global optimization is done by using software-defined networking. The centralized controller keeps track of all of the wavelength entry and exit points and uses routing optimization algorithms and any per wavelength constraints to place them across all of the various sections that make up the network. Examples of constraints include maximum round trip delay and multi-wavelength diversity (when two or more wavelengths need to have separate paths that do not coincide so that they are unlikely to fail simultaneously.)

The packet core is an integral part of the network core. The ingress packets are aggregated on to common links at different points in the network and the role of the network core is to take the packet traffic from the edge, and send it to the designated destination, utilizing the shared optical transmission layer. Since service providers also offer a number of IP services such as Consumer Internet, Business Internet, Connections to other Internet Companies (known as Peering), Private Business Networks (known as Virtual Private Networks), Connections to Cloud Data Centers, etc., the core network also needs to perform the task of bulk packet transport, provided in as service agnostic a manner as possible. Given the different types of IP services that need to be supported, the Multiple Protocol Label Switching (MPLS) is used as a simplified method for packet transport. MPLS uses a small header in front of every packet which allows the core routers on either side to process and forward the packets in a fast and efficient manner. The label portion is an address known to just these two routers (i.e., it is locally significant) allowing re-use of the MPLS label address space on every link. The rest of the header contains information for traffic prioritization (also known as Quality of Service, QoS) used when links become congested.

To manage traffic, a technique called traffic engineering (TE) creates one or more tunnels from every core router to every other core router. Multiple tunnels are used to allow different end-to-end paths to more effectively use core capacity. Inbound packets are mapped on to the tunnel that is appropriate for its destination. Intermediate core routers forward MPLS packets solely based on the tunnel mapping. In case of a link failure, two mechanisms come into play. First, every link has a pre-defined backup re-route path to its neighboring core routers. Second, as the reachability information is propagated across the network, all of the tunnels re-determine their paths so that they no longer expect to use the now failed link. This global repair allows for better global optimization of core capacity. Using software-defined networking, all global traffic information and link state can be accumulated in the SDN controller for even better global optimization. This

allows further tuning of the network. This hybrid approach of using both distributed and centralized control was selected to best balance speed of initial restoration with maximum efficiency. In the near term, a new approach called Segment Routing (or SR), is being introduced. SR allows simplification of the network since multiple control plane protocols can be consolidated. Essentially SR works by allowing each packet to carry path routing information. Thus at the origination point, some or all of the intermediate hops can be pre-determined. Thus, SR supplants the need for separate protocols such as RSVP-TE to execute FRR local repair and LDP for label distribution. This is discussed in more detail in Chapter 13.

The final component of the core is the route reflector. All of the edges that connect to the core need the ability to communicate reachability of service with each other. This is done using the Border Gateway Protocol (or BGP). But how do you let hundreds of edge routers communicate control plane information with each other? By using an aggregation and distribution node known as a route reflector (RR). Every edge router in the network connects to a pair of RRs. For scale, multiple pairs of RR can be deployed for different services or geographic regions. Each of these take in the BGP messages from the edge routers they serve and replicate and distribute the messages to the other edge routers. This offloads the work from the edge and core routers allowing for scale and separation of control and data plane processing.

2.5.10 Service Platforms

For most of the 20th century, the basic capabilities of a network were the customer service, a fixed point to point circuit between two customer locations or the ability to request a temporary narrow band voice connection between two network end points through a switch. Investment in a service platform beyond the core network was largely focused on the OSS, BSS systems and processes to support the delivery and monetization of the network. Business functions included customer ordering, billing, and directory assistance based on a subscription business model. Operational functions included inventory, planning, installation, maintenance and diagnostics. The customer services of transport and switching were an integral part of the network design, expanding the availability of this basic universal service was more attractive than enhancing services that would require investment to upgrade many end points.

As fundamental improvements in transport and computing technology reduced unit cost and increased capabilities of transport and end point devices, enhanced services became feasible. Circuits were replaced by packets, switching replaced by session establishment, and service platforms based on IP protocols emerged to support: simpler and faster connection setup; a range of end point devices that were programmable; increased security of data in transit; interoperability between different services and networks; policies to control network behaviors based on service or customer; multiple types of standard data and media streams; and delivery of rich context to enhance communication.

As technology advanced and architectures converged on IP, customers themselves demanded more than circuits and media sessions, the need for a common platform

became evident from the isolated islands of infrastructure dedicated to a service. The following paragraphs describe some common platform component capabilities used in an all IP network, two examples platforms that integrate common components, and some examples of customer services. The component and capabilities include: directories, public key encryption, signaling protocols, gateways, communication optimizations, and policy rules and enforcement. The platform examples include: the 3GPP IP Multimedia Subsystem (IMS) and a Content Delivery Network (CDN).

While the network was the service, the method for finding an end point was to assign unique numbers to physical end points, provide some level of assurance that the end to end physical connection was secure, have the end user or device find and remember unique numbers they wanted to connect with, and involve the end user to verify the identity of the other side of the connection. As technologies like IP protocols and mobile computing devices became available a better method was required given the diverse range of services expanding number of end points and identities.

Two major classes of technologies support finding and trusting an end point - directories and public key encryption. In modern IP network services, directories are used at multiple levels such as to translate invariant, global unique identifiers to a current IP address, to find the next hop to move a packet destined for a particular IP address toward the destination, and to represent the current devices that are registered and available for a particular type of real-time communications session. Directories may also include rich information about the end point of a connection that assists in the selection and insures the connection is secure. Public key encryption relies on a key pair, one publicly advertised (in directories or other sharing mechanisms) and one held securely by a connection end point. One side of a communicating link between end points, can encrypt information using the public key and the other can decrypt it using the private key. This technique along with a trust hierarchy of directories and hardware methods for securing private keys in devices are used to insure the identity of endpoints and protect private data in transit across public or untrusted networks. For example, the Universal Integrated Circuit Card (UICC) in a mobile phone contains a Subscriber Identification Module (SIM) with keys used to identify and authenticate subscribers.

Another capability in services platform is a standard protocol for connection establishment, referred to as a signaling or control protocol. The signaling protocols enable rapid connection setup, optimization of network resources by deferring allocation of resources until there is agreement between end points or users that communications can begin, setting expectations with end users or devices about how a session will be initiated, negotiating to characteristics that both end points can support, and adjusting communication service algorithms to work acceptably over varying network conditions. Examples of signaling protocols include: SIP (session initiation protocol) that is semantically similar to the traditional caller/called party interaction over a telephone network augmented with rich session descriptor options to support many types of real-time media; XMPP (extensible messaging and presence protocol) a streaming message protocol that has been used for short messaging, presence indications, and to coordinate real-time media connection establishment; and RTCP (real-time media control protocol)

that enables monitoring and feedback on the performance of a real-time connection where packet loss, available bandwidth, and jitter may change and the service adapts use of the network to handle these variations and with minimal impact to user experience.

Another common capability in a service platform are gateways. These are used to maintain compatibility with legacy services as new services are introduced and to limit, enhance, and/or capture information exchanged across network boundaries. Some classic examples of gateways include; a circuit switched telephony gateway that translates control protocols and transcodes media between a circuit switched connection and an IP packet oriented connection; a session border gateway which performs similar functions as the circuit switch gateway and limits the exposed internal end points; and a proxy server that terminates internal or external sessions, forwards information between the two sessions, and restricts, captures, or enhances the information being forwarded.

In an all-IP network, optimizing higher level communications patterns is another platform capability. Some common optimizations include mixing or bridging many media streams into one combined stream that can be sent back to all sources (e.g. mix all real-time audio session in a conference, forward the last N active video session frames to all sources in a video conference), caching frequently accessed files near the edge of a network to reduce capacity required over the wide area network with a content distribution network, and accelerating perceived application performance by constructing a full web page near the sources of data when the page is composed of many small, separate objects, that would take much longer to assemble from a distant location where each session establishment is delayed by end to end protocol exchanges.

To support different services in the network cloud and react in a coordinated way to real-time events requires an ability to adjust algorithms distributed through the network based on a service, specific customer, network condition, etc. These capabilities are referred to generally as policy and includes both the capability to define policy rules and to execute these rules at distributed policy enforcement points.

The context for communications established in advance of an exchange of information can greatly improve the efficiency and effectiveness of communication. An early example of this was the capability to provide caller identification for incoming voice calls to a business and support controls to transfer a connection from a public network user to different location, agent, and/or automated application based on the caller identity. Communication services and platforms continue to expand in their ability to provide a rich context for communication enabled by: mobile access networks; sensors present in mobile hand held and embedded devices; willingness of users to share information on their location, activities, etc.; data collection and analytics used to infer context from prior activity; service provider privacy policies and information controls that allow a user to opt-in to different levels of context data handling.

The 3GPP specified IMS (IP Multimedia Subsystem) is one example of the above components above being used together in a platform. IMS architecture refers to four major layers, service applications, session control and management, access networks, and

end point devices. This layering integrates traditional telecom services as well as provide the platform framework for future multi-media services over IP on mobile phones. IMS uses SIP (Session Initiation Protocol) to negotiate and establish connections, a Home Subscriber Server (HSS) directory, public key encryption with private keys stored in the HSS and UICC/SIM of a connected device, gateways between networks that adapt media, separate media and control, and maintain the security and integrity of the platform.

Significant services and platforms have emerged that do not rely on a complete IMS framework. The session control and management layer is not always needed and/or alternatives for components in this layer may be less costly and complex to deliver a wide variety services. Capabilities such as signaling and identity are frequently tailored to and tightly integrated with a web application, community, and usage scenarios.

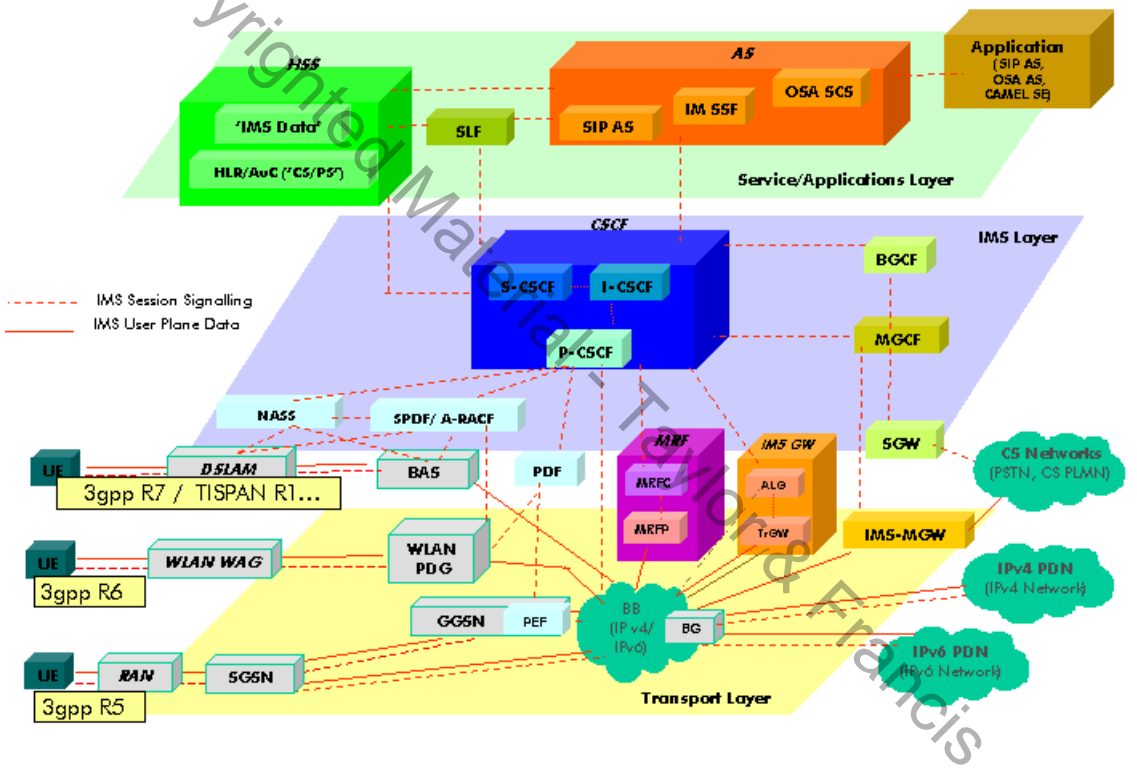

Figure 2.4 - IMS Multimedia Sub-System [4]

A content distribution network (CDN) is another example of a service platform used to optimize the delivery of content that many customers want at the same time in the same general location. A CDN uses components and capabilities such as directories that map a request for content to the closest edge location, streaming media protocols that adapt to network availability and optimize performance under varying conditions.

[4] 3GPP TISPAN

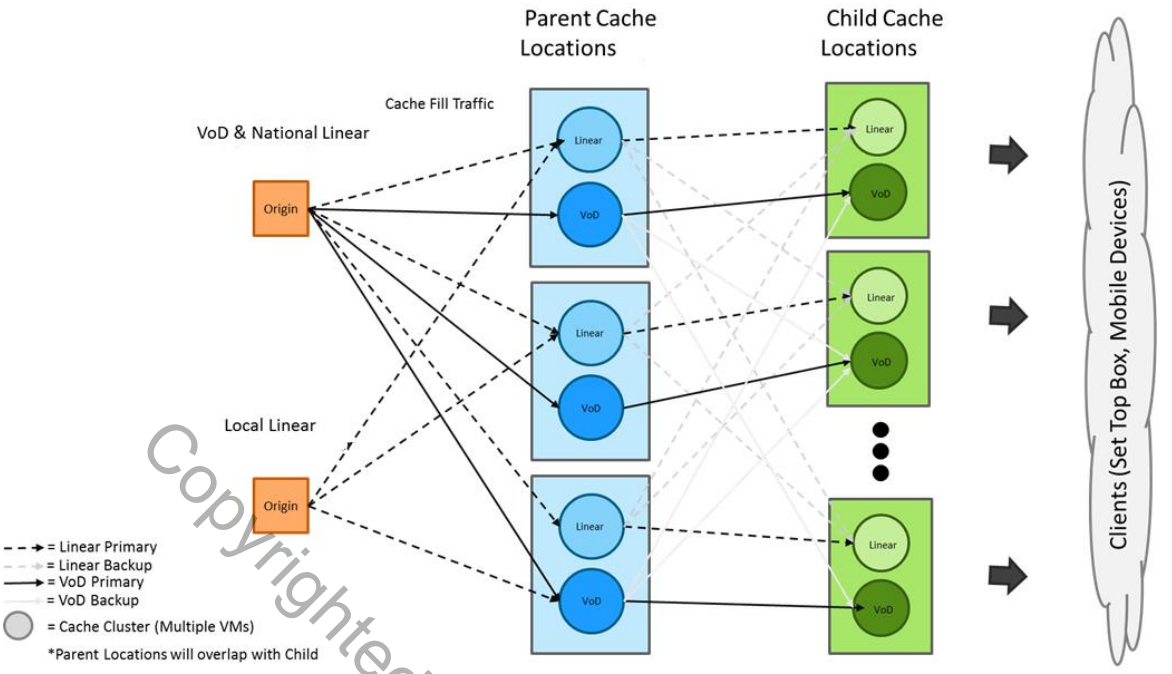

Figure 2.5 - CDN supporting IP video distribution: VoD/linear IPTV

Common examples of business customer services include: internet access; virtual private networks; public and private voice/video networks; and customer dedicated edge configurations that allow connections to a virtual private network at their premise, a public cloud data center, and/or from a wireless device.

Internet access ranges from best effort access over copper wire, fiber, or wireless access point to various levels of managed internet access that guarantee a higher level of service. This may include guarantees on bandwidth, latency, jitter, loss, availability as well as monitoring to detect and protect against security threats. For example, when a distributed denial of service attack is detected adjust routing rules at the ingress points of the attack to drop packets before they are routed to a customer end point.

MPLS Virtual Private Network Service allows a customer to design and control a multi-location, secure, network and establish quotas for how much capacity is devoted to different types of traffic. For example, 40% of the capacity might be dedicated to real-time voice or video to insure an occasional burst of deferrable traffic does not disrupt real-time media connections.

Custom Voice/Video Services are offered to business customers as part of their virtual private network and/or by exposing control of voice/video services from the public telecom network.

Consumer services include the traditional telephone services over copper, fiber, or radio and in the last decade enhanced with broadband IP data that enables many IP applications including those from a large mobile device application ecosystem, IP TV distribution,

and monitoring systems at home and in vehicles. The diverse IP applications enable monetization of network services through multiple business models not limited to network customer subscription.

Chapter 14 describes a wide variety of services, platform components that enable multiple services, and how these are supported in a network cloud.

2.5.11 Network Data and Measurements

In traditional IP networks built with network elements and dedicated to particular services, data and measurements are performed in a relatively static configuration by collecting data from network elements and/or inserting probes at the points necessary to support the requirements of service. For example, one might collect data on real-time media sessions, number of sessions established per unit time, attempts to establish sessions when all capacity is being used, packet loss, jitter, and latency through a session border gateway. By continuously collecting this data it can be used to measure and improve the quality of service and predict when additional transport or switching capacity is needed. Analysis of data in real-time is not critical since the actions one could take based on the data collected involved administrative, planning, and engineering tasks that occur over days or months.

Moving to a network cloud using SDN and NFV creates both challenges to achieving the same level of data collection and measurements as would exist in a traditional IP network and new opportunities to benefit from real-time analytics performed on measurements. The challenges stem from the implementation of a service on flexible infrastructure that is simultaneous supporting multiple services where the mix may change in real-time to real-time collection and analysis across shared, dynamically changing components are opportunities to optimize across multiple services and dimensions. These dimensions can include tradeoffs such as performance of a service, cost to deliver a service, timing of an infrastructure investment, excess capacity to eliminate emergency maintenance, etc. An additional benefit is that tightly coupled to network element service, network, and component measurements are now decoupled and can be done once in a general way within the ONAP platform, reducing the cost, complexity, and time to create services.

Real-time control to optimize the network relies on SDN in the most general sense that all resources must be controllable via software and information is more centralized in near real-time for making decisions. For example, this includes controlling and configuring tunnels, packet flows, load balancers, directories, optical wavelengths mapped to packet layers, etc. and extends all the way to customer services and applications that support interfaces to allow a customer to defer or forecast demands in exchange for better price or performance.

To facilitate appropriate real-time optimization, the capture of network data and executing the measurements, is an integral part of the Network Cloud. The infrastructure to capture, retain and retire network data is described in Chapter 8, while the efficacy of techniques to do the required measurements is discussed in Chapter 16.

2.5.12 Network Operations

The transition of network operations from supporting traditional and IP networks to a Network Cloud is perhaps the largest challenge and key to realize the benefits from the transformation.

The network operational challenges include growing new skills and managing through an interval where there is a mixture of legacy and cloud infrastructure. Significant new software engineering and quality assurance skills are needed to support continuous delivery and integration of new or rapidly changing capabilities, a role commonly referred to as DevOps. Investment to replace current technology infrastructure and refactoring network functions to operate in an automated, closed loop control platform will require operating in a mixed infrastructure for multiple years.

The benefits derived from a successful transformation of operations include: reduced costs; reduced cycle time and increased flexibility to use the network cloud infrastructure supporting new or rapidly changing demand for services; and an increased utilization of infrastructure that is not dedicated but shared with software service when needed.

Operations scenarios that change to realize these benefits include: manual administration, configuration and monitoring of infrastructure or services is highly automated, emergency maintenance is replaced by treating faults as a reduction in capacity that can be addressed during normal, periodic upgrades; planning, testing, and turn-up of new capabilities is routine and highly automated through extensive regression tests delivered with software functions; and organizationally the need for staff aligned to particular services is reduced or eliminated and common technical skills can be shared across all services. Chapter 15 describes the challenges, changes, and approach for an operating a network cloud.

Thinking more broadly than network operations and automating or optimizing a task in the next year to other business functions and creating a culture that embraces change, innovates, and continues to create new value, it is important to consider other types of change. As software has transformed other industries there are examples where major impacts are attributed to a software mentality or discipline that is different than where a telecom provider is today. Two examples of what might sustain leaderships as a network cloud platform provider are: empowering individuals that have the confidence and capability to curate, understand enough to integrate, and invest time to contribute back to peers in a community providing an open source component; and a culture promoting continuous learning, improvement, delivery to a long term platform project emphasizing simple modular designs and unrestricted collaboration with anyone internally or externally that might contribute. Chapter 17 covers many examples of successful operations in a software world and contrasts them with where a telecom provider may be coming from.