OPEN NETWORKING
FOUNDATION

Framework and Architecture for the Application of SDN to Carrier networks

July 18, 2016

ONF TR-534

OpenFlow

ONF Document Type: Technical Recommendation
ONF Document Name: onf2014.183

Disclaimer

THIS SPECIFICATION IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

Any marks and brands contained herein are the property of their respective owners.

Open Networking Foundation
2275 E. Bayshore Road, Suite 103, Palo Alto, CA 94303
www.opennetworking.org

Revision History

Revision	Date	Revision Editor	Changes
1.0	9/6/2016	Shahar Steiff	Release 1.0

Editors: Weiqiang Cheng, China Mobile
Paul Doolan, Coriant GmbH
Dean Cheng - Huawei
Shahar Steiff – PCCW Global

Table of Contents

List of Figures

Abstract

Operators are seeking the benefits promised by SDN to decrease OPEX and CAPEX, while at the same time enabling new applications and revenues via network programmability. The initial ONF focus on SDN has been in the context of data centers (DCs), where the cost-benefit-time tradeoffs were immediately attractive. While there is benefit in extending SDN to the Telco carrier environment, the carrier problem domain introduces additional factors to be considered. Furthermore, even where there are SDN considerations and factors common to the DC and Telco carrier, some of these may have somewhat different emphasis.

This document provides a framework for the specification of SDN in Carrier Networks. This encompasses provision of:

- High level considerations for SDN deployment in Carrier Networks.
- Reference architecture and supporting use-cases/analysis.
- A list of Requirements for SDN deployment in Carrier Networks.

1 Scope

1.1 Scope and Objectives

Carrier SDN focuses on usage of SDN in Carrier Networks, which may offer a wide range of Services supported over a variety of network architectures and technologies. Delivery of an end-to-end service in such networks may involve traversing multiple network domains (e.g., mobile, access, core and data centers) operated by one or more Service Providers, and utilizing resources associated with the supporting infrastructure. The expectation is to achieve a high level of interoperability among the multi-technology, multi-vendor and multi-layer resources of Carrier Networks. SDN can enable binding additional service element (e.g. applications, content) into a complex end-to-end service offering, but it is beyond the scope of this document.

While carrier architectures and their evolution objectives may differ in terms of service mix offered, scalability considerations and other factors, many requirements are uniform across carriers:

- Ability to deliver managed services end-to-end.
- SLA compliance.
- Inter-Carrier interoperability.
- Inter-Vendor interoperability.
- Interoperability/co-existence between SDN and Legacy networks.
- Service Operations and Maintenance capabilities.

This document provides a framework for the specification of SDN in Carrier Networks for the purpose of delivering Network Services.

1.2 Common Terms, and Abbreviations

ACL	Access Control List
A-CPI	Application-controller plane interface
API	Application Interface
BSS	Business Support Systems
Capex	Capital Expenditure
CPE	Customer Premises Equipment
CPI	Controller Plane Interface
DC	Data Center
D-CPI	Data-controller plane interface
E2E	End-to-end
FW	Firewall
FWaaS	Firewall as a Service

HVAC	Heating, Ventilation and Air-Conditioning
IOT	Internet Of Things
LB	Load Balance
JSON	JavaScript Object Notation
LBaaS	Load Balance as a Service
LSO	Lifecycle Service Orchestration
MEF	Metro Ethernet Forum
MPLS	Multi-protocol Label Switching
NaaS	Network as a Service
NAT	Network Address Translation
NBI	Northbound Interface
NE	Network Element
NFV	Network Function Virtualization
NFVI	Network Functions Virtualization Infrastructure
NNI	Network-to-Network Interface
OAM	Operation, Administration and Maintenance
OF	OpenFlow
ONF	Open Networking Foundation
OOB	Out-of-bound
OPEX	Operating Expense
OSS	Operations Support System
QoS	Quality of Service
SC	Service Chain
SDN	Software Defined Networking
SLA	Service Level Agreement
TOR	Top Of the Rack
VAS	Value Added Service
VM	Virtual Machine
VNF	Virtualized Network Functions
VPC	Virtual Private Cloud

VPLS	Virtual Private LAN Service
VPN	Virtual Private Network
VPNaaS	VPN as a Service

1.3 Definitions

Carrier[1]: An organization that owns and administrates a network that provides services.

Operator: An alternative and common name for Carrier; traditionally referring to voice networks.

Service Provider: An organization (Carrier or integrator) that provides network services to the ultimate customers.

Carrier's Network: The network over which that *Carrier* provides connectivity and data transportation.

Platform Domain: One or more devices operated by a single *Carrier*, fulfilling a specific purpose. A *Platform* can be defined by Geography (e.g. "Trans-Atlantic Platform" or "Thames-Valley Platform"), or by Function (e.g. "Access Platform" or "Core Platform"), or by Purpose (e.g. "Managed IP-VPN Services Platform" or "Dedicated Ethernet Platform").

Carrier Domain: May consist of one or more *Platform Domains* – all operated by the same *Carrier*.

Multi-Platform Service: Service that spans across more than one *Platform Domain* – all operated by one single *Carrier*.

Multi-Carrier Service: Service that spans across more than one *Carrier Domain*.

Multi-Domain Service: Service that spans across multiple *Platforms* and/or *Carrier* domains. Both *Multi-Platform Service* and *Multi-Carrier Service* fall under the definition of *Multi-Domain Service.*

[1] The terms of Carrier, Carrier Network, Operator, and Service Provider are for the purpose of this document. They are often interchanged depending on local business jargons.

2 Introduction

This clause discusses some differences between SDN implementation in a Carrier Network and SDN implementation in a DC (Data Centre), as well as possible other applications of SDN.

A Carrier may operate both DCs (which are largely the same as DCs operated by any other organization), Voice (Fixed and Mobile) Networks, and Data (Fixed and Mobile/Wireless) Networks.

2.1 Connectivity and Resiliency

A DC may effectively have non-blocking mesh connectivity between the nodes within the DC, and thus consideration of bandwidth and diversity (e.g. path, technology, policy) constraints may not be a significant Technical, Operational or Commercial issue. However, these constraints are major considerations for Carrier Networks outside of the DC.

> Note - Issues related to connectivity between DC locations are considered external to a DC.

Due to comparatively high cost per bit between nodes in a Carrier Network, the large distances and the effect of external environmental hazards through which the facilities traverse (e.g. damage by third party construction works, submarine cable cuts) issues such as congestion, diversity, latency, resiliency and recovery from facility faults are of high priority.

Communications between controllers and network elements external to a DC are constrained by bandwidth, availability, path diversity and latency factors that are largely absent from control-channel communications within DCs.

2.2 Technology Diversity

A Carrier may operate a DC utilizing a relatively limited number of technologies, vendors and network protocols.

A Carrier Network, on the other hand, has typically included a large number of data plane technologies supporting multiple network protocols. Thus a Carrier Network will need to deal with a combinatorially greater number of technology interactions.

2.3 Network Coverage

The footprint of Carrier Networks may involve large numbers of physical NEs distributed over a multitude of geographically dispersed locations. The need to move large amounts of data over substantial geographic spans implies specialized technology (e.g., DWDM, 40/100G, coherent optics) that will likely remain proprietary for a long while.

2.4 Differentiated Services

The introduction of Managed Differentiated Services, e.g. 5G, will require real-time capabilities of analyzing demand and allocation of resources to address the same.

2.5 Maintenance and OAM

The concentrated nature of a DC implies practices for power, redundancy, maintenance and security that widely differ from the practices needed by Carrier Networks. DC fault detection and isolation may be largely a matter of local monitoring points (power, HVAC, etc.). Failover of service from a faulty VM to another VM is standard practice in a DC environment with limited effect on service performance, if any at all.

Carrier Network Services typically consist of multiple subordinate services that span across multiple network segments and platforms, possibly layered. As a result, fault isolation may be complex. In addition to that - in-service replacement of equipment, particularly if located in remote cabinets or huts with limited capacity for space, power, cooling, and fiber or copper cross-connect, is far more difficult than replacing servers or even TOR switches in DCs. Failover on a Carrier Network will typically result in re-routing of traffic to an alternative path which may exhibit different latency or other performance characteristics. Failover time may vary, depending on protection mechanism. The above may have adverse effect on Service quality (in comparison to a straight-forward failover of service from a faulty VM to another VM in a DC).

2.6 Depreciation

Carrier Networks are typically equipped with a wide variety of devices, provided by multiple vendors, spanning a wide variety of technologies, serving a wide variety of purposes. This equipment will typically depreciate over several years. DCs are typically equipped with a limited variety of equipment types (such as x86 servers, TOR switches), which depreciate rapidly.

> Note 1 – Carrier Network equipment depreciation schedules are often constrained by regulators. This is less likely to be true for DC equipment.
>
> Note 2 – Migration toward NFV will supplant purpose-built equipment to the extent that functionality can be disengaged from specialized hardware.
>
> Note 3 – At any given time a Carrier Network is likely to consist of several generations of equipment and unlikely to be homogenous.

2.7 Services

A DC primarily provides Compute and Storage services, which depend on the type of application being used, regardless of whom the end user is and where it is located.

Traditionally - a Carrier Network provides connectivity services between customer locations. Such services are often agnostic to the application being deployed though they may need to match specific criteria to support the needs of certain applications. Carrier Networks will need to support a very wide range of options and service attributes. The introduction of 5G and other emerging technologies will require faster response of the network to changing demand.

2.8 Migration

The migration from a legacy environment to SDN in Carrier's Networks is usually a brown field upgrade practice in an existing complex environment where service disruption must be minimized or avoided, also ensuring co-existence of SDN and legacy technologies in parallel. A DC may have been built for SDN from the beginning. Yet, minimizing disruption while moving tenant load from one VM in an old environment to a new VM in the new environment, is likely to be a simpler problem. The ONF Migration Work Group has published a document [1]-listing a number of use cases for the Migration into SDN.

2.9 Network Programmability

Carriers often look at SDN beyond its basic concepts of *Centralized-Control* and *Separation of Control Plane from Data Plane*, and consider the broader scope of its *Network Programmability* capabilities, rather than its architecture. Such capabilities allow applications and users to configure and manage network operations according to their requirements. This is typically done indirectly through the Carriers' BSS/OSS platforms rather than a customer directly interacting with the SDN controller. Carrier Networks and the services Carriers offer to their customers, often span across more than one single platform. Thus the use of standards-based protocols and standard-based Information-Models is required in order to effectively program the networks' *end-to-end* behavior. This feature provides tremendous benefit to network operators as well as customers.

The effect of Network Programmability must be well analyzed and studied before real deployment. Agility in Carrier Networks deployment and practices is crucial, and it may be that well defined SDN interface alone might not suffice and the integration must broaden into the OSS and BSS layers.

3 Definition and High-level Requirements for Carrier-Grade SDN

3.1 Carrier Network Overview

Carrier Networks are vital components of regional, national and global infrastructures. Carriers provide a vast array of services to large numbers of retail and wholesale customers including individual users, enterprise customers and governments.

Service Providers hold the role of *network architects*, designing networks and integrating solutions to support the requirements of their service offerings. Carrier Networks are often deploying a wide range of technologies, resulting from infrastructure evolution strategies, that addressed various considerations and constraints related to their operational and business management system environments. Periodical technology and architecture upgrades introduce new best of breed products/vendors that may have emerged. The advent of SDN is one such opportunity.

Typical Carrier Network characteristics encompass:

- Multiservice: the network transports a wide array of services (e.g. fixed and mobile voice, data, media content), addressing a wide area of users (e.g., residential, enterprise, government). Varied consumption of services by different customers leads to varied consumption of network resources with time (e.g., busy hour occurrence) and space (e.g., users' mobility). Because of their scope and importance, services are required to have high agility, availability, reliability and security.
- Multi-vendor: Carrier Networks are built on assets from a variety of suppliers, including vendors of network equipment, OSS/BSS, infrastructure, etc. The main reason for that is that networks are built with intent of using best-in-class (or most-suitable-in-class) products for each segment in a manner promoting competence incentives between vendors.
- Multi-technology: A Carrier Network typically addresses a variety of access and transport technologies spanning different geographic scopes (e.g. metro-access, aggregation, national and international core). This variety serves the multiservice purpose mentioned before, but in many cases it is also the result of technology evolution, legacy equipment persistence, roll-out considerations, infrastructure availability, etc.
- Multi-role: A Carrier Network may play different roles in different service areas or countries. A Carrier can be the incumbent operator in one country or region, offering solid infrastructure availability, while in other countries the same Carrier Network can be based on renting third-party network capacity to complement coverage.
- Multi-Carrier: In certain instances, it takes more than one Carrier to deliver a specific service either because of service design, geographical span or other reasons.

These characteristics typically translate into complex service delivery processes. Service Customization requires individual design and fulfillment resulting in slow adaptation of the network to changing service demands, which then again result in a long time-to-market of new services. The situation is further exacerbated by usage of manual configuration processes, and layer/technology specific network management systems of varying complexity and scale which make the per-service-instance configuration/modification cumbersome and slow.

3.2 Problem Analysis of Current Carrier Networks

Current network architectures and operational models have evolved over decades. To a great extent the networks and the OSS/BSS platforms were designed to provide simple long term services. In recent years Carriers have been subject to increasing pressure from their customers for customized service creation and delivery with very short lead times. This presents a challenge that legacy Carrier Network systems were not designed to cope with. SDN may be part of a solution to some of said issues.

Amongst the problems we observe are the following:

- Multiple technologies, layers and domains result in network resource isolation and fragmentation.

  Carriers in possession of a variety of network resources (e.g. mobile networks, transport networks, IP core network, data centers, etc.) are able to provide end-to-end/cloud-pipe-user services. These resources are located in different domains, employ different technologies and use different management systems. Even simple network operations such as service creation, change or cessation, when manually configured across multiple technologies and management systems with potentially multiple provisioning points/teams, could lead to resource fragmentation and stranding. Complex cross layer/platform coordination is required to enable new services or change existing services. Optimization of a fragmented Carrier Network then becomes very difficult and complex. As a result, resources are utilized conservatively and networks are not operated as efficiently as they might.

- Interoperability issues within and between Carriers.

  Specific Customization for specific Carriers: Network equipment and software vendors often integrate specific software and hardware features, with control plane and forwarding plane tightly coupled, to fulfill the specific dedicated requirements of specific Carriers. In certain scenarios such customization may result in increased CAPEX and OPEX. While in the past customized network gear was common, in recent years Carriers tend to use off-the-shelf gear, but still deploy different software versions and different features and configurations. In the OSS/BSS space Carriers still use tailored solutions, typically a mix of internal development and third party solutions that include different levels of customization. Such customization often becomes a barrier to integration.

  Standardization challenges: The purpose of standardization is to ensure interoperability both between Carriers (when delivering end-to-end service across multiple Carriers) and between different platforms operated by the same Carrier, allowing Carriers to mix and match equipment and management platforms from different vendors. Today – while protocols are well defined and standardized, different service implementation options exist, limiting interoperability between platforms/Carriers. Even a simple service such as a point-to-point Ethernet circuit can be defined by each Carrier using different attributes with inconsistent values assigned to those attributes by different players. Standards are developed in multiple standardization entities that follow different schedules; vendors implement these standards according to different schedules. *The industry is witnessing isolated islands of standardization but lacks end-to-end integration of service delivery standards and lack of a coherent end-to-end Information Modelling capability.* This becomes even more problematic when attempting to replace manual inter-platform configurations with APIs, where lack of a common standard requires tailoring of each and every pair of platforms to share an API. Carriers are also constrained by testing and certification cycles for new equipment and management platforms. Consequently, although communication technology is using standards (maybe too many of them), in most cases, interoperation between different vendors' equipment is still difficult and interconnection between different operators remains challenging, in part as a consequence.

- Lack of flexibility

  The services and resources of legacy network are often tightly-coupled, which results in inability to create new types of services or modify existing services rapidly. This leads to loss of customers and revenue and/or increasing cost of network operations (e.g. duplicate resources/infrastructure).

- Low efficiency of management and operation due to complexity of the network.

  Rapid development of applications and rapid growth in numbers of users leads to rapid growth in the number of nodes and connections the Carrier Network needs to support; the types and volumes of services and the requirements of the customers are growing rapidly, which results more complex in the Carrier's Networks. Traditional networks are usually based on manual and segmented provisioning and operational processes. The new demand from customers presents a challenge to management and operation of networks. Automation is often required in order for Carriers even just to be able to handle the increased workload and complexity of Network Operations.

- Lack of usage of open interfaces

  Providing a flexible open interface is rather complex, especially when the services offered through such open interface transit different network segments (that may use different technologies and may be managed by different OSS) within the Carrier's overall network. This becomes even more complex when parts of such services are provisioned on partner or third-party networks that are not directly managed by the Service Provider. Providing a customized (abstract) view of the network to a customer is complex as such view depends on the products and services offered. This impedes the ability of Carriers to rapidly provide innovative (new) services to their customers.

3.3 Carrier-Grade SDN High-level Requirements

It is well understood that network protocols should be designed so that different parties involved in the delivery and consumption of network resources can communicate with each other, consumers can make use of the resources of providers, and providers can interconnect with each other to provide service. The design must also allow for all the parties involved in the delivery of a service to have the ability to express preferences about which other parties they interact with and express service performance and quality requirements. More generally, network architectures should be designed to enable multi-vendor/multi-Carrier interoperability that enables service providers to offer services to users, as well as enable service provider choice of technology/product, vendor, physical routes etc. Designing for openness, which allows for choice, is also a key aspect of fostering innovation.

While it is possible to create a monolithic proprietary solution that is well designed and partitioned, with well tested internal interfaces, the proof of such an implementation is not apparent to a Carrier except as a result of field experience. Modular systems, with clear demarcations between network layers, are easier to design and validate since they require conformance to standardized interfaces for both the data plane and the control plane. Technology can evolve on either side of the interface independently, not only creating less risk but allowing for independent and optimized technology migration. Modular systems with standardized interfaces provide the Carrier with the ability and flexibility to optimize service delivery to the end-user.

Software-defined networking (SDN) offers the opportunity for resolving many of the challenges enumerated in the previous section via its design for openness. In particular, providing open interfaces that enable the development of software that can control the connectivity provided by a set of network resources and the flow of network traffic through them, along with possible inspection and modification of traffic that may be performed in the network. SDN allows a controller to manage a wide range of data plane resources, and offers the potential to unify and simplify their configuration [2].

In developing SDN for Carrier applications designers may include or exclude certain features, may define certain interfaces, protocols (open or proprietary), in a manner that has a profound influence on operational flexibility and the types of services ultimately delivered to the end user. Carriers must be able to fulfill various customer requirements quickly through modification of, or modifications induced by the application plane. Additionally, the network can condition the application by providing feedback through those interfaces (e.g., congestion notification or packet error rate that could trigger specific actions in the application).

When introducing SDN into Carrier Networks it is the objective that Carriers maintain overall control of their network architectures – control of their own destiny and ability to provide the level and type of service to their customers that they see fit.

When we consider using SDN to address the above-mentioned challenges of Carrier Networks, the need for Carrier-Grade (CG) SDN is obvious. Two aspects of CG SDN should be considered:

1. An SDN enabled Network should inherit the attributes of the original Carrier Network:

Requirement #1*: *High availability

The transport component of traditional Carrier Networks typically requires high availability (e.g. “5 nines”).

Requirement #2*: *High reliability and diversity

Carrier Network services are provided to geographically dispersed customers. To minimize the impact of a disaster and secure the service continuities, not only redundancy of node level, but also that of site- or area- level is required. The architectural design of CG SDN must be considered as logically centralized, but physically distributed controller deployment. In a network with large geographic scale, the protection for the communications infrastructure needs to be considered, and any centralized controller must also support redundancy.

Requirement #3*: *High security

Carriers operate in an environment with many customers/tenants and multiple potential security threat vectors. Security and privacy must be guaranteed. This applies to the interaction with external applications as well.

Requirement #4*: *Manageability and maintainability

Operators need simple and effective tools to operate and maintain the network. Fault monitoring and performance measurement are required to keep the network status and services visible. Automated management and operation processes are required to operate various services to large numbers of customers with complex networks comprising multiple technologies, layers and vendors. Manageability and Maintainability of the network is a prerequisite for on-demand service provisioning and immediate fault analysis and recovery. Inventory management should be also considered.

2. It should be ***future-proof*** and able to address requirements that may not yet be fully known or defined.

Requirement #5*: *End-to-End network resources collaboration

Carrier-grade SDN should be able to coordinate the isolated and fragmented network resources from an entire network perspective such as between mobile and backhaul, between inter- and intra-DC networks, between IP and Optical or between different Carriers involved in the delivery of a service. With this collaboration, Carriers can get the network resources optimized, service capabilities improved, and eliminate manual processes through automation of service lifecycle management.

Requirement #6*: *Network flexibility

There are three aspects where SDN can improve flexibility:

1. Commercial aspects: From a CAPEX aspect, the Carrier Grade SDN should give Carriers the flexibility to choose solutions from different vendors. From an OPEX aspect, Carrier-grade SDN is required to facilitate quick delivery of new services without interruption of

ongoing services. SDN decouples the forwarding and control planes with standardized interfaces, which will provide openness and programmability for the Carrier Networks.

2. Flexible use of existing gear: The lifecycle of equipment can be extended by the re-programming capabilities facilitated by SDN. E.g., aged equipment with lower processing capacity and less throughput rates can be moved toward the network edge and functionally re-programmed, facilitating assets reuse.
3. Increased granularity in network upgrades.

Requirement #7: ***Network intelligence***

Carrier Grade SDN should support very large scale of networks with thousands of network elements and millions of traffic streams. These volumes of traffic and number of streams keep growing year on year, as well as the number and frequency of actions that need to take place on same. Network intelligence such as optimal traffic path selection mechanism per customer and maximizing network utilization must be implemented. Programmability and open APIs can facilitate rapid implementation of such functionality. With the increase in intelligence embedded into the controller, the less human intervention is required. This applies both to automation, service capabilities, application awareness and efficient use of network resources.

Requirement #8: ***Service-aware networking***

Carrier Grade SDN networks should provide the capability to allow end users and customers to directly define new or existing services with specific characteristics such as performance assurance, route etc. An automatic interaction between services and underlying resources should be established without any manual configuration. There is a need for coordination and orchestration mechanisms to convey configuration information between service and underlying resources, as well as reliable consistency check procedures for such an automated provisioning framework.

The services are de-coupled from the specificities of the underlying resources through abstraction. However, it is necessary to ensure that proper mapping exists between service requirements and resources. Coordination among service-related control and resource-related control functions is required, while de-coupled, facilitating differentiated evolution of both.

Requirement #9: ***Network Openness***

The openness of Carrier Grade SDN includes two parts. One is to allow Carriers to customize service orientated interfaces to adapt various requirements from service consumers (e.g. end user, orchestrator, OSS). The other is to facilitate third party functionality integration into the SDN controller.

Requirement #10: It should ***interact with legacy networks and systems*** in a consistent way.

This applies both during the transition period during which services from the Legacy network are gradually migrated to SDN, and during the operation period where certain parts of the legacy network remain in place and must function in tandem with the SDN parts.

4 Carrier-Grade SDN architecture

Carrier Grade SDN should include end-to-end network resources such as supporting infrastructure, mobile, access, core, DC etc. On one hand, Carrier Grade SDN should enable high level of interoperability among multi-type, multi-technology, multi-vendor, multi-layer resources of Carrier SDN networks; On the other hand, it should analyze the requirements of Carrier Grade forwarding, controller, applications and interfaces.

The architecture defined in this document is in compliance with the ONF SDN Architecture [2].

4.1 Types of Open Interfaces in Carrier SDN architecture

a) **Northbound Interface:**

Hierarchical northbound interfaces are provided by the SDN controller to facilitate the flexibility and openness for both service consumers and developers.

b) **Southbound Interface**

Southbound interfaces are provided by the SDN controller to configure and control the data plane, either directly or through an open interface provided by the device OS.

c) **East-Westbound Interface.**

East-Westbound Interfaces are provided by the SDN enabled network to facilitate inter-domain activities in environments where different parts of the network are managed and controlled by different orchestrators (and/or different operators). Such interfaces may also exist in the data-plane layer (physical NNI). The actual implementation of the east-westbound functionality may be handled through existing north and southbound interfaces, however – such interface must abide by the unique operational and business constrains of inter-Carrier and inter-domain functions as detailed in Section 4.2 herewith. East-Westbound Interfaces will typically be between domains complementing each other either in terms of geographical coverage/footprint (extending coverage of a Service Provider beyond the footprint of their own network into the network of a neighbor partnering Service Provider) or in terms of complementing technology (e.g. connecting a Managed-IP-VPN network with a Managed-Cloud platform to allow bundling Cloud services with Connectivity services).

The Carrier Grade SDN controller exposes service orientated interfaces to different service consumers (e.g. end users, orchestrators, OSS) allowing creation of higher level services. To enable the flexibility and openness, the network model and service orientated interfaces can be customized by the Carrier. There are multiple approaches to orchestration of inter-domain and intra-domain services. The details of such approaches are briefly discussed in Section 4.2, Section 9.4 and Annex 1 herewith but are beyond the scope of this document.

Third party functions are also easy to be integrated to the Carrier Grade SDN controller with well defined programming interfaces and possible network programming language.

4.2 Platform Domains, Carrier Domains, Single and Multiple Domains

Depending on the Customer requirement and Service Provider capabilities, the service may be provided by a single Service Provider using their own resources, or span across multiple Carriers and aggregated by the Service Provider into an end to end customer-facing service.

A Service Provider may provide the service using a single network platform (e.g. provide a point to point connection on their MPLS platform) or using a combination of several platforms (e.g. providing Cloud-Application access using their Ethernet Access platform, their MPLS core and their Data-Centre all linked together to provide the end to end service).

An even more complex scenario is where the service spans across multiple Carriers, each delivering their respective part using multiple internal platforms.

We will thus define the following:

Platform Domain: One or more devices operated by a single *Carrier*, fulfilling a specific purpose. A *Platform* can be defined by Geography (e.g. “Trans-Atlantic Platform” or “Thames-Valley Platform”), or by Function (e.g. “Access Platform” or “Core Platform”), or by Purpose (e.g. “Managed IP-VPN Services Platform” or “Dedicated Ethernet Platform”).

Carrier Domain: May consist of one or more *Platform Domains* – all operated by the same *Carrier*.

Multi-Platform Service: Service that spans across more than one *Platform Domain* – all operated by one single *Carrier*.

Multi-Carrier Service: Service that spans across more than one *Carrier Domain*.

Multi-Domain Service: Service that spans across multiple *Platform* and/or *Carrier* domains. Both *Multi-Platform Service* and *Multi-Carrier Service* fall under the definition of *Multi-Domain Service.*

In its simplest form the architecture of a *Platform Domain* will resemble the classic SDN architecture diagrams (Figure 4-1):

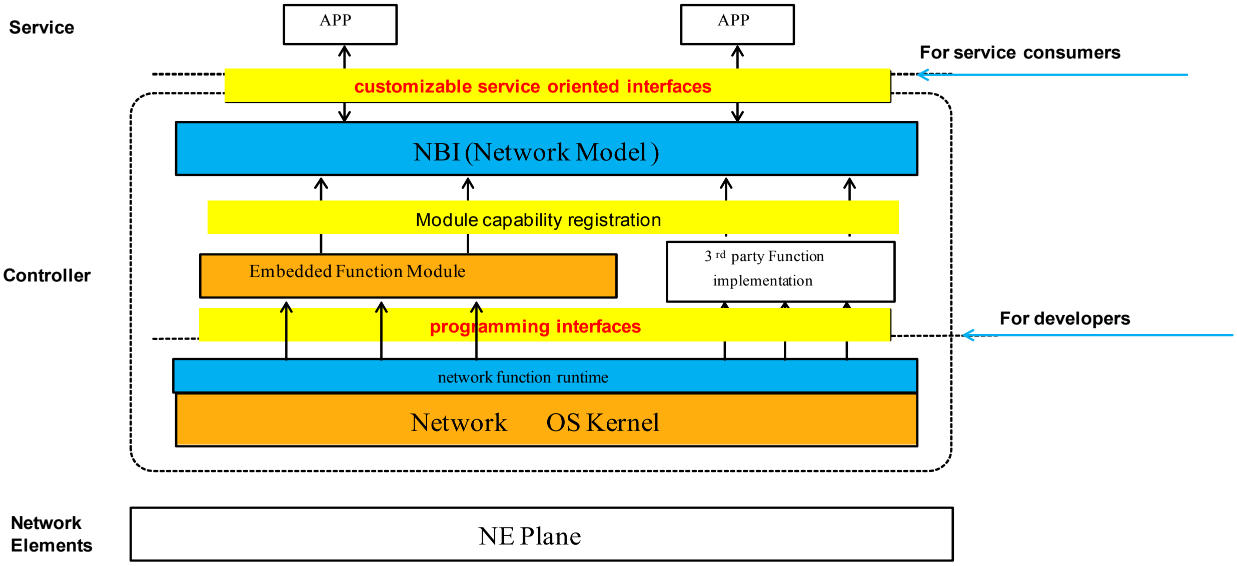

Figure 4-1: Single Network Operator – Single Platform

In the more complex case of a *Multi-Platform Service* we will have multiple controllers, each controlling its platform, and an orchestrator layer (or layers) that orchestrate the internal domains to provide the end to end service (Figure 4-2):

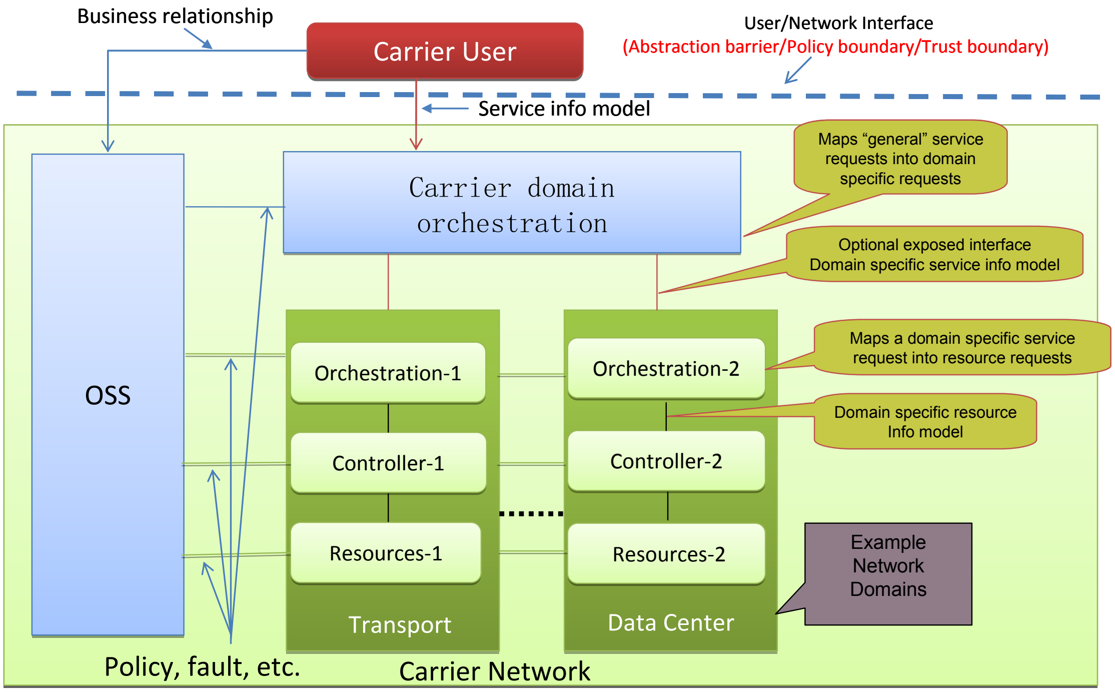
Business relationship
Carrier User
User/Network Interface
(Abstraction barrier/Policy boundary/Trust boundary)
Service info model
Carrier domain orchestration
Maps "general" service requests into domain specific requests
Optional exposed interface Domain specific service info model
OSS
Orchestration-1
Orchestration-2
Maps a domain specific service request into resource requests
Domain specific resource Info model
Controller-1
Controller-2
Resources-1
Resources-2
Example Network Domains
Transport
Data Center
Policy, fault, etc.
Carrier Network

Figure 4-2: Single Network Operator – Multi Platform

Multi Carrier solutions (Figure 4-3) are a bit more complex both from an Operational context and from an Orchestration context.

Operational context: While in a Multi-Domain environment operated by a single Carrier there can be full inter-platform visibility and trust, as well as certain levels of inter-platform controllability (e.g. one platform can view the resources of its neighbor platform and can request activation of such resources directly with its neighbor platform), in a Multi-Carrier environment there is only limited (if any) visibility of one neighbor into its neighbor platforms, thus the inter-controller communications across domain boundaries will be administered through permissions and rules applied by the orchestration/OSS/BSS layers.

In practice – Specifically in the context of Managed Services in a Multi-Carrier environment - Visibility and Control are typically handled either *indirectly* – through the OSS/BSS layers of both Carriers, based on Inter-Carrier operational arrangements those operators may have in place; or *directly* by the controllers on both sides, based on permissions and rules defined by the respective orchestrators/OSS/BSS on both sides. To an extent, the *direct* approach indicated here may evolve to the convergence of the Control-Plane and certain elements of the OSS.

> Note – It is obvious that Inter-Carrier physical connectivity through a shared facility (e.g. NNI) is implemented on the Data-Layer, but configuration and control of the devices on each side of such shared-facility will be administered either directly or indirectly as described above.

Orchestration context: Orchestration provides two major functions:

1. Breaking a complex service request into service components that are recursively supported by domain platforms;
2. Constructing an integrated, end-to-end, solution across multiple platform domains.

There are multiple approaches to handling the first aspect of orchestration defined above, which are beyond the scope of this document and the Carrier environment does not add complexity or special considerations.

Single Carrier orchestrators can achieve end to end visibility through a hierarchical topology of controllers and layers of orchestrators that aggregate information and orchestration capabilities to a single top-level orchestrator. As discussed above, in a Multi-Carrier orchestration scenario there is limited inter-Carrier visibility and control is subject to policies, thus while the top-level orchestrator will be able to directly orchestrate services within its own Carrier Domain. Visibility, Control and orchestration of the elements of service that are provided by neighbor Carriers can be handled through an east-west-bound interface as described in the Operational context above.

> Note - While theoretically Operators could allow visibility and control of their networks to neighbor orchestrators, this would pose scalability and business information privacy issues.

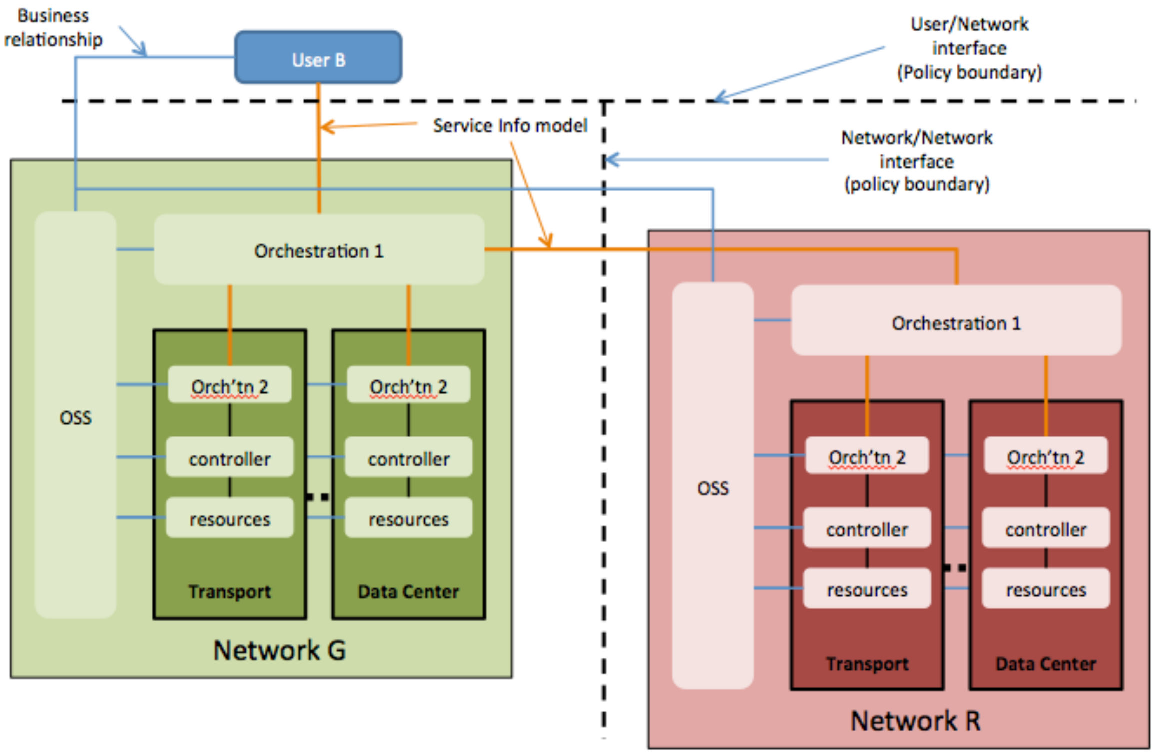

Figure 4-3: Multi Operator – Multi Platform

Note: Orchestration in a Multi-Carrier environment is a relatively new topic and has, to-date, only seen limited attention from the various industry fora. A discussion of some of the current work on Multi-Carrier orchestration can be found in Annex 1 herewith.

5 SDN Based End Customer Service Requests

One of the primary benefits of SDN in the Carrier environment is the ability to automate creation of new services via a programmatic interface to the end customer or application. The benefits to this approach are obvious – it allows customers to receive their services almost instantaneously, rather than having to wait hours, days, or even months. Carriers enjoy earlier revenue from these services, increasing the profitability of their network services. Carriers are also able to offer and deliver new types of services, and service combinations, they were unable to deliver in a manual fashion, thus enable extracting additional revenue from existing resources. This section discusses the categories of customer service requests, and provides examples of how they may be provided in an SDN-based infrastructure.

These service requests fall into one of two broad categories:

Type 1:

- Request expressed in "service semantics"
- Identifies the characteristics of the service (not the resources to provide that service)
    - E.g. "I want a 10GbE link to a 2TB storage server"
    - Such requests can be based on standardized service descriptions developed by industry fora and conveyed through an NBI implementation such as Intent-based requests [3].
- To satisfy the service request the Service Provider will need to translate the Intent-Expressed Service requirements into Data Plane Resource requirements.
    - E.g. new forwarding table entry on a node

Type 2:

- Request using Data Plane Resource requirements.
- E.g. configuration of resources in a virtualized environment.
- To satisfy a service request the Service Provider may need to:
    - Obtain resources from another Carrier and/or
    - Use resources from different (internal) domains (e.g. Transport, Mobile, Data Center)

The methods and processes by which the Orchestrators and controllers interpret those requests and turn them into configuration of resources is beyond the scope of this document, and vary by the type of resources and services involved.

6 Operational aspects

Figure 6-1 illustrates the way many Carrier business-oriented services (for example point-to-point Ethernet services or multipoint IP VPN services) are typically operated today:

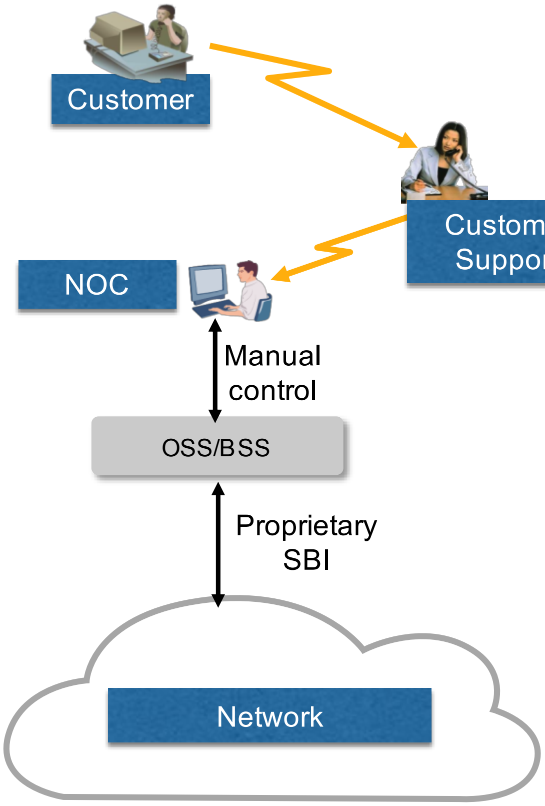

Figure 6-1: Current Carrier Network Operations

In this example customers communicate with the Carrier's customer support center, which then turns the customer requests into service order tickets for the Network Operations Center (NOC). The NOC personnel convert the service order into a set of provisioning or network control commands which are entered in the appropriate OSS/BSS. The OSS/BSS in the figure represents the set of OSSs that are currently deployed in a Carrier's Network and support functions such as Billing and Accounting, Inventory, Network Planning etc. The interface between the OSSs and the network is typically proprietary or perhaps partially standardized (with private enhancements). Because of the manual intervention and rigorous internal processes required to complete these operations, the time between the customer requesting a service and fulfillment of such request can be relatively long (days or weeks) even if the network resources (e.g. fiber access to the customer's location) already exist. This may have a negative impact on customer satisfaction and on Carriers' revenue streams.

SDN has the promise of providing Carriers with standards-based tools to streamline and automate many of these processes as illustrated in Figure 6-2. Automation enables the option of avoiding human intervention when requesting new services or changing existing services. The scope of control that a customer is allowed must be negotiated between the customer and the

Carrier, these constraints are enforced by the policy enforcement function shown in Figure 6-2 below. Typically such requests are made through a customer service portal that is interacting with the controller either directly or through the OSS/BSS platforms.

Figure 6-2 shows how operations can be improved in an SDN environment:

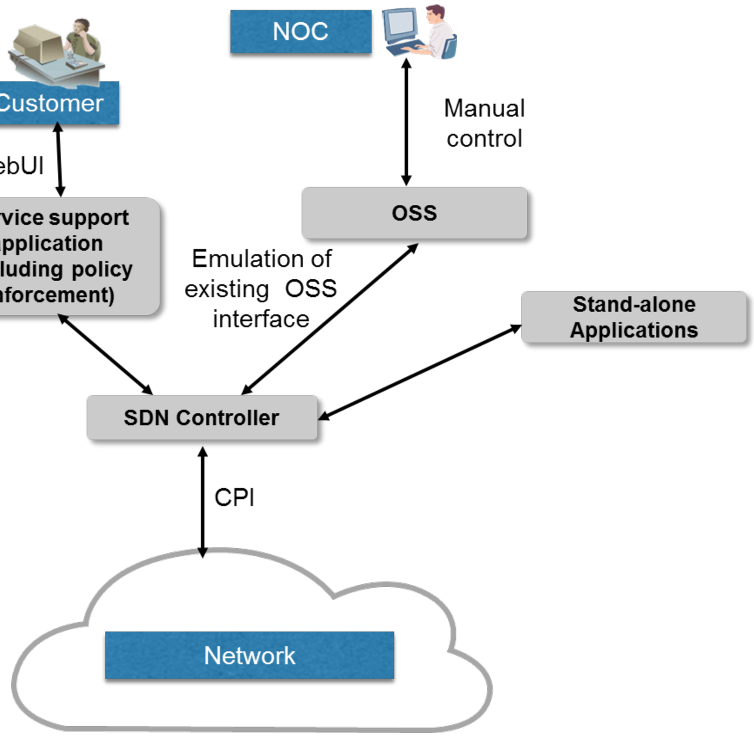
NOC
Customer
Manual control
WebUI
Service support application (including policy enforcement)
OSS
Emulation of existing OSS interface
Stand-alone Applications
SDN Controller
CPI
Network

Figure 6-2: SDN-based Carrier Network Operations

To fully realize the potential operational simplifications made available through Network Programmability (SDN), some level of automatic network optimization is required. These functions may be present in the SDN controller, the existing OSS's, in a standalone application (as shown in Figure 6-2) or in any combination of these. These automated planning and provisioning tools can also use network performance measurements, such as utilization and SLA conformance, to make automatic changes to the network to improve the network's performance.

Reducing the manual operations has a number of benefits. It allows the Carrier to respond faster to customer requests; allows customers to have direct control of their network services; is more responsive to changes in the network's performance; reduces human errors.

7 NFV and its relationship with Carrier Grade SDN

7.1 NFV in Carrier's Networks

Network Functions Virtualization, or NFV, is a technology developed by ETSI NFV ISG and other standards organizations to resolve some of the challenges discussed in previous chapters. The primary vision of NFV is similar to that of SDN: to separate hardware and software on network equipment, which are then consolidated onto servers (compute), storage-units (storage) and switches (connectivity) based on industry standards.

These two different approaches are highly complementary to each other. The focus of SDN is to allow software and applications to control traffic flow in networks, while the focus of NFV is to virtualize some network functions over standards based hardware so resources can be shared and utilized by software entities.

This issue is discussed in depth in several documents issued by both ONF and ETSI, some of which are listed herewith. It is beyond the scope of this document to go into the definitions of NFV.

- ETSI - Introductive White Paper (published 2012 [4])
- ETSI - NFV Architecture (published 2014 [5])
- ETSI -Usage of SDN in NFV (published 2015 [6])
- Relationship of SDN and NFV (ONF TR-518 published October 2015 [7])

7.2 High-level Requirements on interactions between Carrier Grade SDN and NFV

Both SDN and NFV are emergent technologies that are in initial deployment in Carriers' networks. During the design and deployment of Carrier Grade SDN, Carriers must consider the NFV factor and leverage the benefits of each technology during integration. Note that within a Carrier Grade network, an SDN domain and an NFV domain may belong to separate administrative entities; e.g., the NFV domain may be operated by a third party but provide virtualized resources to the Carrier Grade network's operator who owns the SDN operation. Also, there may be more than one NFV domain within a single SDN-enabled Carrier Network. This section specifies some high-level requirements for Carrier Grade SDN with NFV factor.

The ONF SDN architecture allows an SDN controller to manage a wide range of data plane resources. In SDN-enabled Carrier Networks, some resources may be owned and administrated by one or more NFV domains.

The following requirements are valid for both hybrid environments where the network resources are a mix of legacy and NFV and for environments that are all NFV based:

> ***Requirement #11:*** Carrier Grade SDN must be capable of operating in networks where some network functions are virtualized via NFV domains.
>
> ***Requirement #12:*** A Carrier Grade SDN Controller must be able to manage both virtualized and non-virtualized network functions owned by NFV domains.

ONF has defined the SDN controller as capable of handling generic resources, some of which may be NFV virtual network functions (VNFs) or NFV network services (NSs).
Requirement #13: Management and Discovery of Network functions should allow a Carrier Grade SDN controller to discover the network functions that have been instantiated through NFV.

Requirement #14: An SDN application should be capable of collecting performance related information concerning the networking resource usage, if provided, by NFV domain.

Requirement #15: A Carrier Grade SDN operation must not violate security measures and policies of an NFV domain deployed in the same network but owned and operated by another organization.

Requirement #16: The management functions in an SDN-enabled Carrier Network must be capable of managing both virtualized and non-virtualized resources in NVF domain.

The following requirements are valid for environments where the network resources are all NFV based:

Requirement #17: A Carrier Grade SDN controller should be capable of (re-)creating and removing a VNF instance owned by a given NFV domain.

Requirement #18: An SDN application should be capable of managing and verifying the configuration of the elements that virtualize the hardware resources, if provided, by NFV domain.

8 Integration of Carrier Networks and Data Center Networks

8.1 Service Requirements

Carrier Networks (refer to Figure 8-1) may consist of several domains, such as Metro, IP backbone and Data Center ("DC") networks. From the Carriers' perspective, each domain has its unique requirements and challenges, some of which are:

- Balancing traffic in Metro networks
- Traffic steering in backbone networks
- Multi-tenant and service chaining in DC networks

From the customers' perspective, the expectation is a self-managed end-to-end service that can be dynamically provided on-demand. By utilizing SDN, Carriers can provide such network services to customers with higher efficiency, performance and quality. This is generally known as NaaS (Network as a Service).

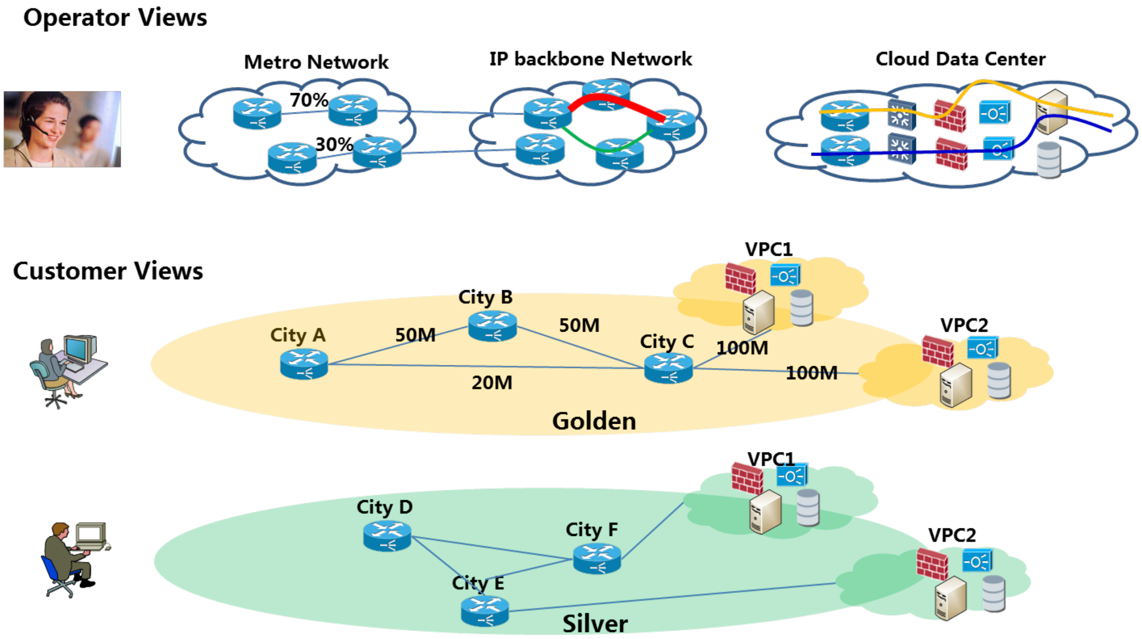

Figure 8-1: Service Views

In Carrier's Networks, the requirements to provide NaaS are as follows:

Requirement #19 VPC and Service Chaining in one or multiple DCs

a. Virtual Private Cloud (VPC): Carriers should allow their customers to define and manage their cloud-based network, compute and storage in real time, through an SDN controller. The network services provisioned typically include IP addresses, network subnets, ACL, QOS, FWaaS, LBaaS, and VPNaaS. Note that the VPC service may exist within a single DC or across multiple DCs.

b. Service Chaining (SC): Carriers should allow their customers to define the service chaining for their north-south or east-west traffic through an SDN controller. For example, in public and private cloud, the service chaining functions may include NAT/FW/LB/VPN. Note that SC service may exist within a single DC or across multiple DCs.

c. Carriers should support heterogeneous operation, i.e., while some DCs are SDN enabled, other DCs may use legacy technology.

Requirement #20 Flexible VPN in Managed network environments.

Flexible VPN: Carriers should allow customers to define the connectivity and required bandwidth of their VPNs across the Carriers' Managed network environments (such as MPLS, VPLS, and Carrier Ethernet etc.) using SDN enabled on-line portals. The Carriers should be able to provision such VPN services in near-real time.

Requirement #21 End to End (E2E) VPC & Flexible VPN & Service Chain

Carriers should allow E2E services initiated by their customers to span over multiple domains (such as backbone network, DCs). Such services may include connectivity, bandwidth, SLA, lifecycle management, VAS, SC, etc. Such E2E services may be bundled into one combined product offering.

Requirement #22 Traffic optimization in Managed network environments.

Carriers should deploy logically centralized SDN controllers in order to distribute the network load evenly, improve overall bandwidth utilization, and guarantee end-to-end QoS for customers' applications.

With the anticipated increase in traffic volumes, some network elements/areas may become congested. With current networking technologies, networks are deployed such that each and every network device is designed and configured to have the capability to handle the peak level traffic, though it is unlikely that all network elements will experience peak traffic at the same time. Apparently this approach is ineffective and inefficient causing waste of network resources. SDN's centralized architecture allows customers and their applications to share network resources through planning, coordination, and virtualization, maximize the network capacity with optimization and greatly enhances the operational efficiency.

A single data center provides one or more network services as described above. A Carrier's network usually connects with one or more data centers, where each data center can be owned by the same organization as the Carrier's network, but can also be owned by others such as another Carrier, an enterprise, etc. To provide integrated network services, data centers may also be interconnected.

8.2 Carrier Networks based SDN Architecture

An integrated Carrier's Network, consisting of multiple segments, may be centrally controlled by an SDN Application Platform (as described in Figure 8-2 herewith), but it is more likely to be segregated into separate domains, and each domain is controlled by its own management tool (which may be SDN based or legacy). The integration and orchestration of those segregated domains through a centralized tool, or a federation of distributed orchestrators, currently remains as a challenge due to lack of standards.

Separation of Control-plane from Data-Plane, when it comes to a multi-domain environment, may require an evolved approach and high levels of integration.

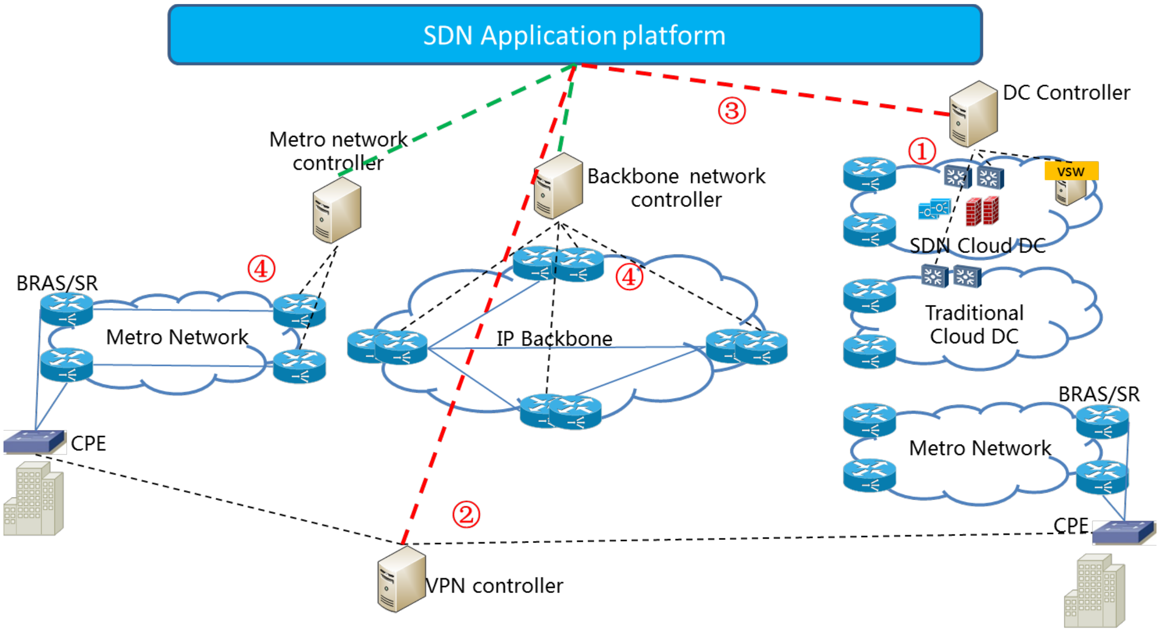

Figure 8-2 : Sample Network Architecture

8.3 Migration Considerations

When deploying SDN in data centers and transport networks, we should consider the integration and co-existence of new platforms with legacy elements of the network. That includes network devices as well as their associated management and support systems. Overlay solutions and software update solutions are typically the preferred methods, considering a gradual migration strategy.

When deploying SDN in an existing network, Carriers will likely prefer to gradually migrate services from the legacy network to the SDN enabled network, rather than migrate all services at once.

This topic is further discussed in Section 9.3.3 herewith.

9 Impact of SDN on Carrier Networks

The introduction of SDN in Carrier Networks presents a number of implications from operation, organization, and business points of view [8]. Some complementary impacts are detailed herewith.

9.1 From SDN-enabled Data-Centre to Carrier Network Programmability

Existing Carrier Transport services are typically *fixed-lines* based on *manual service delivery* with lead times that may span weeks and months and service commitments of months or more. The demand, however, is shifting towards *on-demand services allowing assured (e.g. QoS managed) transport* between nodes. The Carrier industry has become a commodity, a utility, providing the pipelines through which third party content is transmitted using best-effort protocols. In order to overcome the best-effort nature of the public internet Carriers strive to deliver assured services. Current assured transport services are primarily provisioned manually thus unable to effectively deliver short-term on-demand services, thus unable to meet customer demand. In this chapter we will demonstrate that SDN is well positioned to enable a new eco-system based on assured, on-demand services.

While the classical data-centre centric SDN paradigm focuses on *separation of control plane from data plane*, and *centralized control*, the Carrier-Network paradigm shifts to ***network-programmability***. Rather than using software to control how a device forwards packets, a Carrier will seek ways to ***use software to manage the entire network***. Most Carriers would be looking for software, open-source or proprietary, to enhance their network in several aspects:

- Automation of service configuration and activation currently performed manually, leading to *Faster Service Delivery* and *Simplified Processes*.
- Introduce the ability to deliver *on-demand, dynamic*, features and services that can *enhance existing Carrier Transport related products* and *add new products* to the portfolio.
- *Improved Situational-Awareness* and *faster fault/problem resolution*. Allowing automatic reconfiguration of network resources based on changing traffic patterns or network conditions.
- Enable *Platform/Domain-Agnostic network management*.

Most Carriers would be looking for software that is able to perform platform-agnostic abstraction of the network infrastructure, using open or proprietary protocols to configure ports and forward packets, and using software to manage network resources and their utilization. Taking a step higher than defining a controller that is directly configuring the network device, equipment vendor may need to introduce an Abstraction-Virtualization layer below the controller. This layer has two functions: Northbound Abstraction: interfaces with the device, builds an inventory of nodes and the possible virtual connections between them and represents them in a standards-based manner. Southbound Virtualization: uses the device OS layer (or an open-source solution) to configure the network devices and build virtual services as requested by the controller. The controller sits above the Abstraction/Virtualization layer, matches the service requirement with the inventory of virtual connections and sends a request to the Abstraction/Virtualization layer to instantiate a virtual connection. The Abstraction/Virtualization layer then configures the network devices using their specific device's management protocol (e.g. Open-Flow/Netconf/YANG/JSON). This approach is illustrated in Figure 9-1.

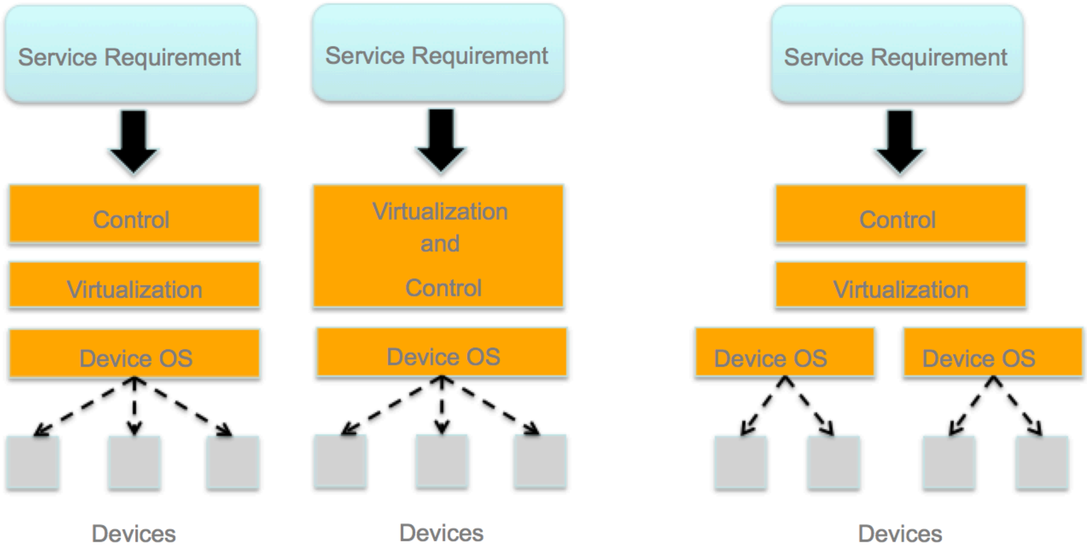

Figure 9-1: Typical Carrier-Grade Architecture

9.2 Programming Responsibility

There are three broad categories of actors we can consider as developers of SDN software:

- Customers
- Carriers
- Vendors

Each of these actors has different domains of expertise and the level of that expertise will vary on a case by case basis. Each of the actors will also have different levels of programming expertise.

SDN applications are pieces of code that use the A-CPI and/or D-CPI. Given the appropriate specification, tool chain and access, any of the above parties can write and deploy an SDN application. In principle - this can be done by any application or user of the network. In practice *policy* will limit who is allowed to program what.

With the caveat that these are generalizations we offer the following observations:

- End customers of Carriers will not be allowed to program NEs at any level. However – Customers and Carriers alike may have interest and the ability to develop applications that access the network via an A-CPI.
- Some Carriers will have the ability or business desire to develop SDN controllers.
- The vendor community (which includes Network Equipment vendors, Solution integrators and others) may also have the ability to develop applications at all levels of the SDN stack.

- Core infrastructure/equipment providers and System Integrators will be the most likely actors to develop software to support the open interfaces required for SDN.
- The Carrier-Grade Orchestrator will likely be integrated into the OSS/BSS platforms either as an integral part or as an add-on to an existing platform.
- The use of Open Source in Operational systems and in Reference Implementations is left to the discretion of the stakeholders – considering the potential benefits and risks. Open Source based Reference Implementations may support validation of components developed in software and can serve as an effective tool in the acceleration of standards' development and validation.

9.3 Discussion of CAPEX/OPEX and Revenue in SDN-enabled Carrier Networks

One of the major "selling points" for SDN in the Data-Centre environment is Cost-Savings through migration from legacy to programmable white-label gear. The promise of CAPEX savings in a Carrier Network may not be as dramatic for several reasons:

9.3.1 Suitability of white-label gear for Carrier-Grade deployment

There are certain areas in the Carrier Network where white-label devices can effectively replace traditional proprietary gear. CPE is a good example. SDN is a key enabler in the application of programmable white-label gear.

In other areas of the network, such as the Core, considerations such as port-density, resiliency, power consumption and others may pose challenges to white-label gear.

9.3.2 Cost of Gear vs. Cost of Network

While in data-centers equipment cost is a significant element in the annual balance sheet, and a moderate reduction in equipment cost will yield a significant impact on CAPEX, in most Carrier Networks equipment cost typically only comes third, after *cost of Human Resource* and the *cost of the transport media* (submarine cables, optical fibers, and copper, including cost of deployment). Cost of equipment may still amount to a large figure, but being a smaller element overall - even a significant reduction in equipment cost will typically only yield a smaller reduction in overall CAPEX.

The cost of orchestration and control infrastructure should also be taken into account. While use of lower cost gear (e.g. white-label) may have a positive effect on overall cost, the associated orchestration and control infrastructure that is required to operate it needs to be considered as well.

The effect of use of SDN on transport cost is thus both direct (through CAPEX reduction) and indirect (through improvement of network utilization and potential introduction of new services).

9.3.3 Cost of migration

Migration from legacy network to programmable network can not be performed overnight. Typically, a Carrier will need to go through some of the steps listed below, or more:

- Design and build of new system in parallel to legacy (technical).

- Training existing staff. Hiring new (mostly expensive - IT) staff (HR).
- Gradual migration of customers to new platforms (operations).
- MarCom (Marketing Communications). While there is promise for new revenue – the potential customers must be made aware of it through marketing campaigns (business)

The costs associated with each of those stages vary case by case.

9.3.4 OPEX Considerations

Migration to a Programmable Network has both positive and negative implications on Operational Costs:

- Staff – Shift of *low-cost skills* (provisioning, operations, and field technicians) to *expensive skills* (System Architects, Programmers).
- *Faster Fault Isolation and Repair* yielding *Shorter Downtime* yielding *Reduction in SLA Penalties*.
- *Improved Network Utilization* yielding *Slower Growth of Network Costs*.

9.3.5 New Revenue Opportunities

Using programmable infrastructure Carriers will be able to offer new services that can be billed by the minute or by the bit. Content that is currently transported over the public internet can be migrated to managed, assured, networks.

This will create a new eco-system where network performance can be managed and applications and content that take advantage of such network capabilities can now be developed and offered.

This represents an opportunity to extract new revenue streams from existing network infrastructure, an opportunity for application developers to offer differentiated performance levels (e.g. video resolution and refresh-rate) through guaranteed bandwidth offering, and an opportunity for content providers to offer differentiated packages based on guaranteed QoS.

- New products that can be realized through on-demand service activation (e.g. QoS managed Video).
- Billable enhancements to existing products (e.g. QoS/Security on demand, bandwidth increase on demand).
- Cross-Platform inter-Carrier services on-demand (e.g. mobile to transport to cloud) in one go – penetrate new markets. The service is delivered across multiple platforms/segments and the revenues can be split between the stakeholders.

9.4 Challenges Faced during Deployment of SDN in Carrier Networks

Several challenges currently exist when considering deployment of SDN in Carrier Networks to enable Network Programmability.

9.4.1 No Multi-Domain Orchestration yet Exists

Programmability, to date, is typically limited to a single device or a single platform. No multi-domain orchestration yet exists (though several such attempts are underway as described in Annex 1). As a result – manual configuration is still required in topologies involving multiple domains/platforms/Carriers.

9.4.2 Service Spans across Multiple Platforms and Multiple Carriers

Network controllers/orchestrators assume centralized control, where a controller/orchestrator maintains the state of nodes at all times so it:

- Knows “everything about everyone” that is within its control-domain (abstracted views included).
- Is able to manage and change status of all objects within its control-domain.
- May use sub-controllers in a hierarchical manner to abstract and control sub-domains.

These assumptions work well in confined environments that are operated by a single Carrier using a single platform, but they break when the service spans across multiple platforms and multiple Carriers, which most services actually require.

9.4.3 Scale Challenges

The ability of a controller to manage state of domains that include a large number of objects may deteriorate with growth in size:

- Number of objects - database management, memory constraints
- Real-time processing of large volumes of information (administrative, operational).
- Field tests show propagation of change-of-state information of a single object in large, distributed, networks may take minutes. Too slow to be considered “real time” or “on demand”.

Such challenges are not unique to Carrier Networks and may be experienced by other segments of the market too.

9.4.4 Distance Challenges

The ability of a controller to effectively communicate and control networks that span across large geographies may be challenged. A logically-centralized controller might not be able to overcome challenges such as:

- Latency caused by distance between centralized controller and remote objects. Delayed communications between controller and nodes.
- Connectivity options between centralized controller and remote objects may lack diversity due to limited/costly infrastructure. OOB (out of band) control may not be available.

9.4.5 Inter-Domain Challenges

When services span across more than one administrative/operational domain:

- Neighbor domains may not offer visibility/controllability of state of its objects to neighbor controllers/orchestrators that are operated by a different operator.
- Neighbor domains do not have hierarchy. They are players in an equal-level playing field. This breaks the assumption of logically centralized (thus hierarchical) control.

9.5 The Federated Multi-Domain Orchestrator

As presented in Section 9.4 above, the assumption of a single logically centralized controller breaks under the multi-domain/Carrier environment. For Carriers to fully adopt SDN and take full advantage of network programmability, a new approach should be adopted where each domain uses a logically centralized SDN controller, but the end to end orchestration is performed cooperatively by the controllers using an appropriate federated multi-domain orientated protocol.

Inter-Carrier operations require a standardized Information-Model and standardized APIs. In an environment where different domains may be operated using different controllers/orchestrators, the use of an *industry-wide common Information-Model* and *industry wide standardized east-west-bound APIs* is imperative.

Requirement #23 The Federated Multi-Domain Orchestrator should:

- Know its neighbors exist and may have limited knowledge of their capabilities, but does not maintain state of their nodes or capabilities except those nodes involved in existing services that span through that neighbor.

- Communicate with its neighbors through a standardized east-west-bound API. Such communications may include queries about capabilities, requests for activation or modification or termination of service, status queries about existing services.

To date - SDN turns manual silos into automated silos, but Carrier Networks are made of multiple silos, and inter-Carrier services traverse multiple silos. SDN needs to evolve its inter-silo capabilities in order to properly address the Carrier Network needs.

10 Annex 1 – Multi-Domain/Multi-Carrier Orchestration

Based on the discussions in previous chapters, it is evident that:

1) An Orchestrator maintains the state of services within its platform domain. It might not necessarily maintain the state of all services in other domain platforms that may take part in delivery of an end-to-end service. It must maintain the end-to-end state of existing services that span beyond its own platform domain. An orchestrator may maintain the state of additional reference points along the end-to-end path of existing services based on operational agreements between the various stakeholders involved in delivery of the end-to-end service.

2) A Hierarchical Orchestration approach that assumes a single-top-level orchestrator that has visibility of the entire mesh of networks, and is capable of configuring services across domains – is impractical. It suffers both scalability and trust constraints.

3) In order to orchestrate the delivery of services across multiple domains, in an environment where an Orchestrator does not have end to end visibility of all possible connectivity options and permutations, we must develop a Federated Orchestration framework based on an approach that allows progressive refinement of a request.

The MEF LSO (Lifecycle Service Orchestration) Reference Architecture [9] and the 5GEx [10] project multi-domain proposal (Figure A1-1) are examples (possibly the only two examples of such an approach at the time this document is written) that allow multi-domain inter-Carrier orchestration.

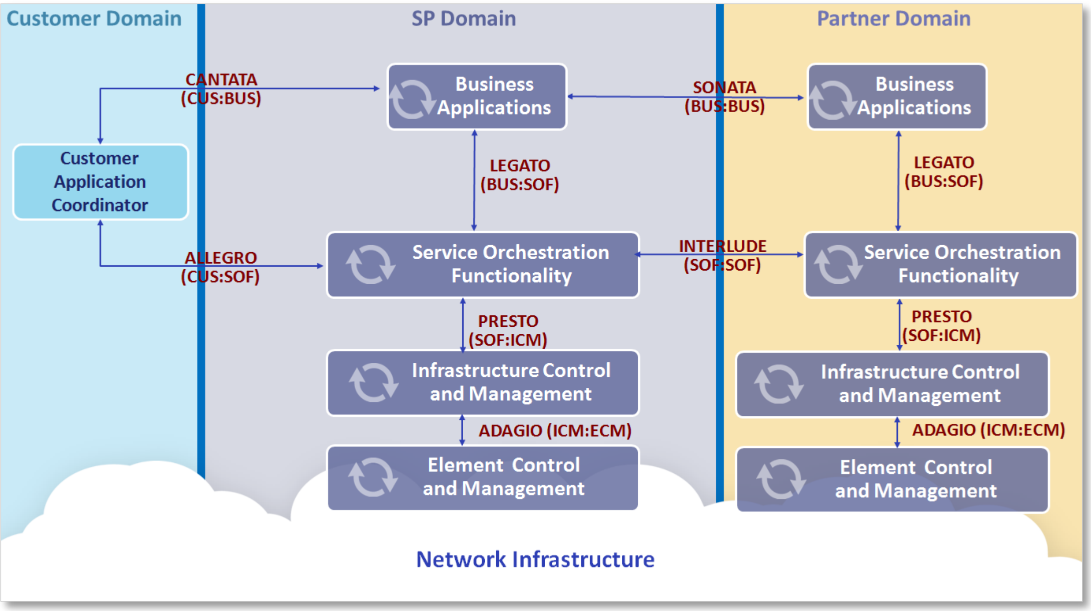

Figure A1-1: MEF LSO Reference Architecture

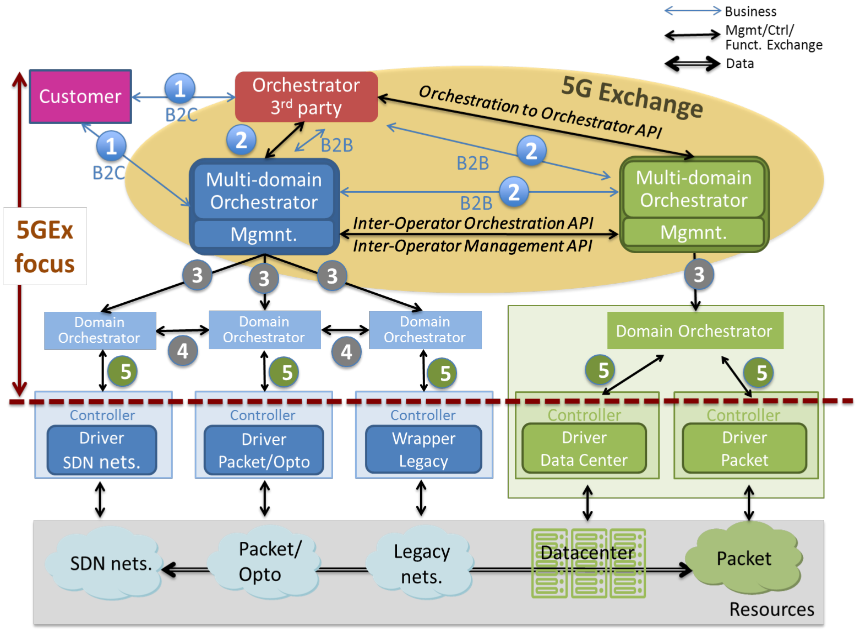

Figure A1-2: 5GEx Reference Architecture

Cooperation (Figure A1-2) between Carriers takes place at the higher level through the Inter-Carrier orchestration API (2) that exchanges information, functions and control (“Sonata” and “Interlude” in LSO, and “B2B” and “Inter-Operator-Orchestration API” in 5GEx). These interfaces serve for the Business-to-Business and Operations-to-Operations relations between Carriers to complement the Business-to-Customer relations. An optional third party Orchestrator can be integrated in certain scenarios (this is only envisaged in the 5GEx project). Domain orchestrators (3&4 in 5GEx, “Presto” in LSO) and controllers (5 in 5GEx and “Adagio” in LSO) operate during the orchestration, control, and enforcement of domain policies required for multi-domain orchestration. This approach allows for a clear demarcation between the inter-domain elements and the intra-domain elements, while still ensuring the flexibility to handle both and keeping local infrastructure details confidential and hidden from neighbors. The multi-domain orchestrator is in charge of abstracting the underlying infrastructure before it announces what utility and functions the operator is capable of to its neighbouring Carriers. Using such an inter-working architecture for multi-domain orchestration will realize use-cases that are nowadays hard to tackle due to the interactions of multiple heterogeneous actors and technologies.

Existing ONF work can be mapped into those frameworks. For example, OpenFlow sits well in the LSO “Adagio” interface, T-API is a possible implementation of LSO “Presto”, and the East-West API discussion group within ONF discuss LSO “Sonata” and “Interlude” implementations.

A similar mapping can be done with 5GEx architecture. For example, interface 3 can be assimilated as I-CPI interface, interface 5 could be implemented as T-API, and interfaces 2 and 4 could leverage on the East-West API.

11 References

[1] ONF, Migration Use Cases and Methods, available at https://www.opennetworking.org/images/stories/downloads/sdn-resources/use-cases/Migration-WG-Use-Cases.pdf

[2] ONF, SDN Architecture Issue 1.1, available at https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR-521_SDN_Architecture_issue_1.1.pdf

[3] ONF Intent NBI - Principles and Definition available at https://login.opennetworking.org/bin/c5i?mid=4&rid=7&gid=0&k1=1556&k2=2&k3=9&tid=1468847594

[4] Network Functions Virtualization – Introduction White Paper, ETSI, October 2012.

[5] Network Functions Virtualization (NFV) - Architectural Framework, ETSI, December, 2014.

[6] Report on SDN Usage in NFV Architectural Framework (ETSI GS NFV-EVE 005 V1.1.1) http://www.etsi.org/deliver/etsi_gs/NFV-EVE/001_099/005/01.01.01_60/gs_nfv-eve005v010101p.pdf

[7] Relationship of SDN and NFV, ONF TR-518, October 2015

[8] L.M. Contreras, P. Doolan, H. Lønsethagen, D.R.López, “Operation, organization and business challenges for network operators in the context of SDN and NFV”, in Elsevier Computer Networks, Volume 92, pp. 211-217, 2015.

[9] MEF-55 LSO reference architecture http://www.mef.net/Assets/Technical_Specifications/PDF/MEF_55.pdf

[10] 5GEx project home page http://www.5gex.eu

LIST OF CONTRIBUTORS

Weiqiang Cheng – China Mobile

Dean Cheng – Huawei

Shahar Steiff – PCCW Global

Li Chen – China Mobile

Luis M. Contreras – Telefonica

Nicolai Leymann – Deutsche Telekom

Vishnu Shukla – Verizon

Feng Wang – China Telecom

Malcolm Betts – ZTE

Paul Doolan – Coriant

Andy Malis – Huawei

Eve Verma – Nokia

Dave Hood – Ericsson

Sibylle Schaller – NEC

Manuel Paul - Deutsche Telekom