Recent Developments in Service Function Chaining (SFC) and Network Slicing in Backhaul and Metro Networks in Support of 5G

Adrian Farrel*, *Managing Director, Old Dog Consulting*

* Tel: (0044) 1978 860944, e-mail: adrian@olddog.co.uk

ABSTRACT

In this invited paper we examine the role of Service Function Chaining (SFC) and Network Slicing in providing support for emerging 5G services as data is carried across backhaul and metro networks.

The way in which backhaul and metro networks are managed and operated affects the ability of those networks to support a wide range of services that will be enabled by 5G connectivity. Those services will place exacting demands on the bandwidth, delay bounds, and jitter delivered to users, but may vary significantly over relatively short time periods. SFC and Network Slicing offer opportunities for software-driven coordination of network resources to marshal them so as to best deliver the network behaviour that will support the demands of the services.

We present the latest proposals towards standardisation of SFC and Network Slicing in the Internet Engineering Task Force (IETF) and discuss how these concepts are applied in optical networks through the Metro-Haul project.

Keywords: Service Function Chaining (SFC), Network Slicing, Network Functions Virtualization (NFV), 5G, Software Defined Networking (SDN), Service Orchestration.

1. INTRODUCTION

5G is a term applied to cover wireless technologies that will deliver growth in the number of clients supported and a radical increase in performance with bandwidths of up to 20 gigabits per second. 5G technologies will be deployed my mobile telephone companies and will be used by human users (through their telephones and computers) and by massive networks of devices for Internet of Things (IoT) connectivity. Machines using the 5G networks may be at fixed locations or may be highly mobile.

Many use cases have been proposed for 5G networking that leverage the increased bandwidth and quality capabilities, and utilise the support for large numbers of mobile end nodes. Many of these use cases involve the delivery of end-user services that will require the deployment of special capabilities (such as storage and compute) within the network. Other use cases place demands on the network to deliver data according to quality demands that extend from basic throughput to guarantees of low latency, bounds on jitter, and resilience to network events.

While bodies such as the 3GPP focus on the wireless and access networks, less attention has been paid to how the traffic for 5G applications is transported across metro and core networks in a way that will support and enable the delivery of new end-to-end services. Various tools will be critical in enabling transport for 5G transport for 5G services including Software Defined Networking (SDN), Service Function Chaining (SFC), Network Functions Virtualization (NFV), and Network Slicing.

This paper briefly introduces these concepts, explains what work has been started within the Internet Engineering Task Force (IETF), and describes the applicability of these tools to optical metro networks with special reference to the Metro-Haul project. The paper concludes by suggesting areas of further work.

1.1 Overview of Concepts

Software Defined Networking (SDN): The ability of software applications to program individual network devices dynamically and therefore control the behaviour of the network as a whole. [1]

Network Functions Virtualization (NFV): The use of standard IT virtualisation technologies to consolidate many network equipment types onto industry standard high volume servers, switches, and storage, that can be located in Data Centres, Network Nodes, and in end user equipment. [2]

Service Function Chaining (SFC): The definition and instantiation of an ordered set of network service functions and the subsequent steering of traffic through those functions. [3]

Network Slicing: An approach to network operations that builds on the concept of network abstraction to provide programmability, flexibility, and modularity. [4]

----

2. FORWARDING PLANE APPROACHES TO SERVICE FUNCTION CHAINING

In SFC, traffic flows made up of data packets must be steered between service functions instantiated at different places in the network. This is addressed as a two-layer networking problem where transport connectivity (such as IP tunnels, MPLS LSPs, or pseudowires) is use to build pipes between "Service Function Forwarders" (SFFs) in the overlay SFC network, and an SFC encapsulation is used to indicate to which chain of service functions the packet has been assigned and which specific function should be exercised next.

Three SFC encapsulations have been proposed within the IETF:

2.1 Network Service Header (NSH)

The IETF's SFC working group has defined the NSH [5] as a layer-independent encapsulation for SFC. The NSH may be seen as a "shim" header that is inserted between the transport encapsulation and the header of the end-to-end user packet. The NSH carries an identifier of the chain that the packet is traversing, an indication of the point on the chain that has been reached (i.e., the next function to be executed), a TTL field that prevents packet loops, and a container for metadata.

The NSH is agnostic of the type of payload packet (it may carry IP, MPLS, Ethernet, or anything) and uses a "next protocol" field so that service functions know how to process the payload. Similarly, the NSH can be carried by any transport protocol as packets are routed across the network.

Support for the NSH may require the development of new hardware capabilities as well as extensions to existing software realizations of network functions.

2.2 MPLS Logical Representation of the NSH

Multiprotocol Label Switching (MPLS) is a widely deployed forwarding technology that uses labels placed in a packet in a label stack to identify the forwarding actions to be taken at each hop through a network. Actions may include swapping or popping the labels as well, as using the labels to determine the next hop for forwarding the packet. Labels may also be used to establish the context under which the packet is forwarded. [6]

New work in the IETF's MPLS working group provides a mechanism to achieve SFC in an MPLS network by means of a logical representation of the NSH in an MPLS label stack. [7] An approach is also provided to install metadata specific to data flows or service chains and to reference it from within data packets.

This technique has the potential to work with existing MPLS forwarding equipment such that SFC features can be more quickly deployed into existing networks using existing equipment.

2.3 Segment Routing (SR)

An old paradigm in packet forwarding – source routing – is seeing a renaissance in the form of Segment Routing (SR) under the care of the IETF's SPRING working group. [8] SR places a list of identifiers in a packet header and uses these sequentially to determine the next routing hop or processing function to be applied to the packet. This list is applied to packets when they enter the network so that the routing state is carried in the packets rather than being distributed through the network.

From its inception, SR has included applicability to SFC amongst its objectives. A service function chain may be encoded as a list service functions and their locations. This list can be applied as the user packets as an encapsulation so that all of the information in order to deliver a packet to the service functions is carried along with the packet.

3. PATH SELECTION AND NETWORK CONTROL FOR SFC

As well as forwarding packets to deliver them to a series of service functions, an SFC system needs to include some other key features. It must provision service function instances within the network, instantiating them on servers or virtual machines, and configuring them to function correctly; it must discover and announce the locations of service functions so that packets can be routed to them; and it must plan and configure service chains so that the correct sequence of functions is executed with consideration of load-balancing and end-to-end delivery characteristics.

3.1 Instantiating and Configuring Service Instances

Instantiation of service functions is a classic NFV orchestration issue. That is, a service function is a software component that is loaded, configured, and executed on a virtual machine hosted on a server.

This type of orchestration requires planning both for CPU load and network load: the service function must be instantiated on a server that can support it; sufficient service function instances must be created to support the demands of the traffic (assuming adequate load balancing across the instances); and the service functions must be present where they are needed in the network so that path lengths are not increased unreasonably by steering traffic to the service functions.

It is this last point that makes placement of service function instances more than just a hypervisor operation: awareness of traffic demands, SFC needs, network topology, and the status of network resources are all important.

3.2 Service Function Discovery and Advertisement

Once service functions have been instantiated and configured, they must be made available for use. In an SFC network, this involves binding each of them to an SFF (usually a process of discovery) and advertising their location so that packets on a specific SFC can be routed between SFFs for delivery to the service functions, and so that service function chains can be planned.

A mechanism for advertising the reachability of service function instances using the BGP routing protocol is being developed by the IETF's BESS working group. [9]

3.3 Planning Service Function Chains

The planning of service function chains can be simple or complex depending on the network features available. A given traffic flow needs to be passed through a sequence of service functions according to the service being delivered and the policies of the network operator; the ordering of the service functions is critical, and the choice of service function instances will depend on current and predicted loads as well as the location of those instances in relation to each other and considering the end points of the traffic flow; the planning process may be made to produce more optimal chains by enhancing it to allow the creation of additional service function instances.

From many aspects, the SFC planning operation is similar to traffic engineering path computation and so may be performed by a suitably enhanced Path Computation Element (PCE). [10]

3.4 Programming SFC Network Components

Once a service function chain has been planned and the necessary service functions have been instantiated, the network must be programmed to forward packets along the chain. In all cases, a Classifier function at the ingress to the network must be instructed how to map packets to a specific chain: this may be a management operation using, for example, a YANG model; or it may use extensions to the PCE communications Protocol (PCEP) such as those suggested in the IETF's PCE working group. [11]

Once the Classifier has been programmed it can start encapsulating packets to traverse the service function chain. This may involve adding an NSH, applying MPLS labels, or adding an SR instruction list. The Classifier can then forward packets towards the first SFF either as a direct operation (such as IP forwarding) or by tunnelling them across the underlay network.

Note that, if the SR approach is not used, the SFFs must also be programmed to know how to forward SFC packets to the locally attached service function instances and then onwards to the next SFF. This information can reach the SFFs from a central controller's southbound interface, or may use a protocol such as BGP [9].

4. NETWORK SLICING AS A RESOURCE MANAGEMENT TOOL

Network slicing is a mechanism for partitioning network resources for specific uses. This may reserve capacity for certain traffic flows (as in traffic engineering), or it may serve to keep traffic from different users separate as is the case with virtual routing and forwarding (VRF) and layer three virtual private networks (L3VPN).

In the context of the support for different services, network slicing may be a useful resource management tool. For example, it can be utilised to construct a virtual network that can be operated by a service provider's customer as though it was a real and private network. Alternatively, a network operator may use network slicing to reserve resources and dedicate them to the support of an application, service, or customer. Thus, a network slice might be constructed to guarantee bandwidth, low latency, and resiliency for a class of service or a specific user – a feature that is seeing a lot of attention in the context of 5G networking.

The IETF has researched various aspects of network slicing to attempt to determine what work may be applicable within its scope. This has led to several "Birds of a Feather" (BoF) meetings that have discussed potential work. The most recent of these meetings (termed the COMS BoF) investigated Common Operations and Management of Slices as set out in [12]. Other proposals have looked for a narrower scope and aim to re-use existing architecture and protocol tools [4].

5. USE CASES IMPACTING THE OPTICAL METRO NETWORK

Although it is possible to conceive of components like amplifiers and opto-electronic converters as chainable network functions within an optical network (such that an optical path must be routed through a sequence of these functions to allow end-to-end delivery), this is not SFC as normally envisioned. That is, SFC is not a process used to apply functions to a path (that is more usually considered as traffic engineering), but is intended to apply those functions to the data (i.e., packets) that travels along a path.

However, many use cases for SFC impact the optical metro network. [13] For example, a large content distribution may require the selection of CDN sources, the routing of traffic through security functions, and the delivery of the traffic to the intended destination through a choice of access gateways. The path chosen must consider the availability of bandwidth between the possible CDN sources, the service functions, and destination gateways, and that may require the establishment of new optical paths or the modification of existing flexible optical paths.

6. NETWORK SLICING IN A CONVERGED METRO NETWORK

Network slicing may be a particularly effective tool in a metro network especially when it is called upon to deliver high-feature services for packet demands such as those arising from 5G [14]. For example, in order to provide a guarantee of high-bandwidth low-latency connectivity for varying demands across a number of sites (such as might happen in a metro network connecting 5G edge nodes supporting different 5G cells, and various data centres and connectivity to the wider Internet). In such a case, a slice of the metro network can be used to reserve the resources necessary to guarantee service delivery and allow separate management of the parts of the network that are supporting the service.

7. OPPORTUNITIES FOR FURTHER WORK

The Metro-Haul project [15] aims to design and build a smart optical metro infrastructure able to support traffic originating from heterogeneous 5G access networks, addressing the anticipated capacity increase and its specific characteristics, e.g., mobility, low latency, low jitter etc. The project has set five targets:

1. 100 times more 5G capacity supported over the same optical fibre infrastructure
2. 10 times less energy consumption
3. Latency-aware metro network in which latency-sensitive slices are handled at the metro edge ensuring the metro network adds no additional latency
4. End-to-end SDN-based management framework enabling fast configuration time to set up or reconfigure services: specifically 1 minute for simple network path set-up, 10 minutes for full installation of a new VNF, and 1 hour for setting up a new virtual network
5. Reduction in CAPEX of a factor of 10, plus a reduction in OPEX of at least 20%.

REFERENCES

[1] E. Haleplidis *et al.*: Software-Defined Networking (SDN): Layers and Architecture Terminology, *IETF*, RFC 7426, Jan. 2015, <https://www.rfc-editor.org/info/rfc7426>.

[2] Network Functions Virtualisation – An Introduction, Benefits, Enablers, Challenges & Call for Action: Issue 3, *ETSI*, Oct. 2014, <https://portal.etsi.org/NFV/NFV_White_Paper.pdf>.

[3] J. Halpern and C. Pignataro: Service Function Chaining (SFC) Architecture, *IETF*, RFC 7665, Oct. 2015, <https://www.rfc-editor.org/info/rfc7665>.

[4] D. King and Y. Lee: Applicability of Abstraction and Control of Traffic Engineered Networks (ACTN) to Network Slicing, *IETF*, draft-king-teas-applicability-actn-slicing-02, Jan. 2018, work in progress.

[5] P. Quinn *et al.*: Network Service Header (NSH), *IETF*, RFC 8300, Jan. 2018, <https://www.rfc-editor.org/info/rfc8300>.

[6] E. Rosen, A. Viswanathan, and R. Callon: Multiprotocol Label Switching Architecture, *IETF*, RFC 3031, Jan. 2001, <https://www.rfc-editor.org/info/rfc3031>.

[7] A. Farrel, S. Bryant, and J. Drake: An MPLS-Based Forwarding Plane for Service Function Chaining, *IETF*, draft-ietf-mpls-sfc-00, Mar. 2018, work in progress.

[8] S. Previdi *et al.*: Source Packet Routing in Networking (SPRING) Problem Statement and Requirements, *IETF*, RFC 7855, May 2016, <https://www.rfc-editor.org/info/rfc7855>.

[9] A. Farrel *et al.*: BGP Control Plane for NSH SFC, *IETF*, draft-ietf-bess-nsh-bgp-control-plane-03, Mar. 2018, work in progress.

[10] A. Farrel, J.-P. Vasseur, and G. Ash: A Path Computation Element (PCE)-Based Architecture, *IETF*, RFC 4655, Aug. 2006, <https://www.rfc-editor.org/info/rfc4655>.

[11] D. Dhody, A. Farrel, and Z. Li: PCEP Extension for Flow Specification, *IETF*, draft-ietf-pce-pcep-flowspec-00, Mar. 2018, work in progress.

[12] L. Geng *et al.*: Problem Statement of Common Operation and Management of Network Slicing, *IETF*, draft-geng-coms-problem-statement-04, Mar. 2018, work in progress.

[13] A. Hmaity, F. Musumeci, M. Tornatore, and A. Pattavina; Protection strategies for virtual network functions placement and service chains provisioning, *Wiley Journal*, Sep.2017.

[14] R. Casellas, R. Martínez, R. Vilalta, and R. Muñoz: Towards 5G slicing in Metropolitan Networks: Integrating TSDN in a NFV context, *Wireless World Research Forum #39*, Oct. 2017.

[15] Metro-Haul Project Overview, https://metro-haul.eu/