Network Working Group   I. Bryskin
Internet-Draft   Huawei Technologies
Intended status: Informational   X. Liu
Expires: December 31, 2017   Jabil
                                                       J. Guichard
                                                            Y. Lee
                                                Huawei Technologies
                                                      June 29, 2017

                  Use Cases for SF Aware Topology Models
              draft-bryskin-teas-use-cases-sf-aware-topo-model-00

Abstract

   This document describes some use cases that benefit from the network
   topology models that are service and network function aware.

Status of This Memo

Copyright Notice

the Trust Legal Provisions and are provided without warranty as
described in the Simplified BSD License.

Table of Contents

1. Introduction

Normally network connectivity services are discussed as a means to
interconnect various abstract or physical network topological
elements, such as ports, link termination points and nodes
[I-D.ietf-teas-yang-te-topo] [I-D.ietf-teas-yang-te]. However, the
connectivity services, strictly speaking, interconnect not the
network topology elements per-se, rather, located on/associated with
the various network and service functions [RFC7498] [RFC7665]. In
many scenarios it is beneficial to decouple the service/network
functions from the network topology elements hosting them, describe
them in some unambiguous and identifiable way (so that it would be
possible, for example, to auto-discover on the network topology
service/network functions with identical or similar functionality and
characteristics) and engineer the connectivity between the service/
network functions, rather than between their current topological
locations. The purpose of this document is to describe some use
cases that could benefit from such an approach.

2. **Exporting SF/NF Information to Network Clients and Other Network SDN Controllers**

    In the context of Service Function Chain (SFC) orchestration one existing problem is that there is no way to formally describe a Service or Network Function in a standard way (recognizable/ understood by a third party) as a resource of a network topology node.

    One implication of this is that there is no way for the orchestrator to give a network client even a ball-park idea as to which network's SFs/NFs are available for the client's use/control and where they are located in the network even in terms of abstract topologies/virtual networks configured and managed specifically for the client. Consequently the client has no say on how the SFCs provided for the client by the network should be set up and managed (which SFs are to be used and how they should be chained together, optimized, manipulated, protected, etc.).

    Likewise, there is no way for the orchestrator to export SF/NF information to other network controllers. The SFC orchestrator may serve, for example, a higher level controller (such as Network Slicing Orchestrator), with the latter wanting at least some level of control as to which SFs/NFs it wants on its SFCs and how the Service Function Paths (SFPs) are to be routed and provisioned, especially, if it uses services of more than one SFC orchestrator.

    The issue of exporting of SF/NF information could be addressed by defining a model, in which formally described/recognizable SF/NF instances are presented as topological elements, for example, hosted by (associated with) TE, L3 or L2 topology nodes. The model could describe whether, how and at what costs the SFs/NFs hosted by a given node could be chained together, how these intra-node SFCs could be connected to the node's Service Function Forwarders (SFFs, entities dealing with SFC NSHs and metadata), and how the SFFs could be connected to the node's Tunnel and Link Termination Points (TTPs and LTPs) to chain the intra-node SFCs across the network topology.

3. **Flat End-to-end SFCs Managed on Multi-domain Networks**

    SFCs may span multiple administrative domains, each of which controlled by a separate SFC controller. The usual solution for such q scenario is the Hierarchical SFCs (H-SFCs), in which the higher level orchestrator controls only SFs located on domain border nodes. Said higher level SFs are chained together into higher level SFCs via lower level (intra-domain) SFCs provisioned and controlled independently by respective domain controllers. The decision as to which higher level SFCs are connected to which lower level SFCs is

driven by packet re- classification every time the packet enters a given domain. Said packet re-classification is very time-consuming operation. Furthermore, the independent nature of higher and lower level SFC control is prone to configuration errors, which may lead to long lasting loops and congestions. It is highly desirable to be able to set up and manage SFCs spanning multiple domains in a flat way as far as the data plane is concerned (i.e. with a single packet classification at the ingress into the multi-domain network but without re-classifications on domain ingress nodes).

One way to achieve this is to have the domain controllers expose SF/ NF- aware topologies, and have the higher level orchestrator operate on the network-wide topology, the product of merging of the topologies catered by the domain controllers. This is similar in spirit to setting up, coordinating and managing the transport connectivity (TE tunnels) on a multi-domain multi-vendor transport network.

4. Managing SFCs with TE Constrains

Some SFCs require per SFC link/element and end-to-end TE constrains (bandwidth, delay/jitter, fate sharing/diversity. etc.). Said constraints could be ensured via carrying SFPs inside overlays that are traffic engineered with the constrains in mind. A good analogy would be orchestrating delay constrained L3 VPNs. One way to support such L3 VPNs is to carry MPLS LSPs interconnecting per-VPN VRFs inside delay constrained TE tunnels interconnecting the PEs hosting the VRFs.

Planning, computing and provisioning of TE overlays to constrain arbitrary SFCs, especially those that span multiple administrative domains with each domain controlled by a separate controller, is a very difficult challenge. Currently it is addressed by pre-provisioning on the network of multiple TE tunnels with various TE characteristics, and "nailing down" SFs/NFs to "strategic" locations (e.g. nodes terminating many of such tunnels) in a hope that an adequate set of tunnels could be found to carry the SFP of a given TE-constrained SFC. Such an approach is especially awkward in the case when some or all of the SFs/NFs are VNFs (i.e. could be instantiated at multiple network locations).

SF/NF-aware TE topology model in combination with TE tunnel model will allow for the network orchestrator (or a client controller) to compute, set up and manipulate the TE overlays in the form of TE tunnel chains. Said chains could be duel-optimized compromising on optimal SF/NF locations with optimal TE tunnels interconnecting them. The TE tunnel chains (carrying multiple similarly constrained SFPs)

could be adequately constrained both at individual TE tunnel level
and at the chain end-to-end level.

5. SFC Protection and Load Balancing

Currently the combination of TE topology & tunnel models offers to a
network controller various capabilities to recover an individual TE
tunnel from network failures occurred on one or more network links or
transit nodes on the TE paths taken by the TE tunnel's connection(s).
However, there is no simple way to recover a TE tunnel from a failure
affecting its source or destination node. SF/NF-aware TE topology
model can decouple the association of a given SF/NF with its location
on the network topology by presenting multiple, identifiable as
mutually substitutable SFs/NFs hosted by different TE topology nodes.
So, for example, if it is detected that a given TE tunnel destination
node is malfunctioning or has gone out of service, the TE tunnel
could be re- routed to terminate on a different node hosting
functionally the same SFs/NFs as ones hosted by the failed node.
This is in line with the ACTN edge migration and function mobility
concepts. It is important to note that the described strategy works
much better for the stateless SFs/NFs. This is because getting the
alternative stateful SFs/MFs into the same respective states as the
current (i.e. active, affected by failure) are is a very difficult
challenge.

At the SFC level the SF/NF-aware TE topology model can offer SFC
dynamic restoration capabilities against failed/malfunctioning SFs/
NFs by identifying and provisioning detours to a TE tunnel chain, so
that it starts carrying the SFC's SFPs towards healthy SFs/NFs that
are functionally the same as the failed ones. Furthermore, multiple
parallel TE tunnel chains could be pre-provisioned for the purpose of
SFC load balancing and end-to-end protection. In the latter case
said parallel TE tunnel chains could be placed to be sufficiently
disjoint from each other.

6. Network Clock Synchronization

Many current and future network applications (including 5g and IoT
applications) require very accurate time services (PTP level, ns
resolution). One way to implement the adequate network clock
synchronization for such services is via describing network clocks as
NFs on an NF-aware TE topology optimized to have best possible delay
variation characteristics. Because such a topology will contain
delay/delay variation metrics of topology links and node cross-
connects, as well as costs in terms of delay/delay variation of
connecting clocks to hosting them node link and tunnel termination
points, it will be possible to dynamically select and provision bi-
directional time-constrained deterministic paths or trees connecting

clocks (e.g. grand master and boundary clocks) for the purpose of exchange of clock synchronization information. Note that network clock aware TE topologies separately provided by domain controllers will enable multi-domain network orchestrator to set up and manipulate the clock synchronization paths/trees spanning multiple network domains.

7. Client - Provider Network Slicing Interface

3GPP defines network slice as "a set of network functions and the resources for these network functions which are arranged and configured, forming a complete logical network to meet certain network characteristics.". Network slice could be also defined as a logical partition of a provider's network that is owned and managed by a tenant. SF/NF-aware TE topology model has a potential to support a very important interface between network slicing clients and providers because, on the one hand, the model can describe holistically and hierarchically the client's requirements and preferences with respect to a network slice functional, topological and traffic engineering aspects, as well as of the degree of resource separation/ sharing between the slices, thus allowing for the client (up to agreed upon extent) to dynamically (re-)configure the slice or (re-)schedule said (re-)configurations in time, while, on the other hand, allowing for the provider to convey to the client the slice's operational state information and telemetry the client has expressed interest in.

8. Dynamic Assignment of Regenerators for L0 Services

On large optical networks some of provided to their clients L0 services could not be provisioned as single Och trails, rather, as chains of such trails interconnected via regenerators, such as 3R regenerators. Current practice of the provisioning of such services requires configuration of explicit paths (EROs) describing identity and location of regenerators to be used. A solution is highly desirable that could:

o Identify such services based, for example, on optical impairment computations;

o Assign adequate for the services regenerators dynamically out of the regenerators that are grouped together in pools and strategically scattered over the network topology nodes;

o Compute and provision supporting the services chains of optical trails interconnected via so selected regenerators, optimizing the chains to use minimal number of regenerators, their optimal

      locations, as well as optimality of optical paths interconnecting
      them;

   o Ensure recovery of such chains from any failures that could happen
      on links, nodes or regenerators along the chain path.

   NF-aware TE topology model (in this case L1 NF-aware L0 topology
   model) is just the model that could provide a network controller (or
   even a client controller operating on abstract NF-aware topologies
   provided by the network) to realize described above computations and
   orchestrate the service provisioning and network failure recovery
   operations.

9. Dynamic Assignment of OAM Functions for L1 Services

   OAM functionality is normally managed by configuring and manipulating
   TCM/MEP functions on network ports terminating connections or their
   segments over which OAM operations, such as performance monitoring,
   are required to be performed. In some layer networks (e.g.
   Ethernet) said TCMs/MEPs could be configured on any network ports.
   In others (e.g. OTN/ODUk) the TCMs/MEPs could be configured on some
   (but not all network ports) due to the fact that the OAM
   functionality (i.e. recognizing and processing of OAM messages,
   supporting OAM protocols and FSMs) requires in these layer networks
   certain support in the data plane, which is not available on all
   network nodes. This makes TCMs/MEPs good candidates to be modeled as
   NFs. This also makes TCM/MEP aware topology model a good basis for
   placing dynamically an ODUk connection to pass through optimal OAM
   locations without mandating the client to specify said locations
   explicitly.

10. SFC Abstraction and Scaling

   SF/NF-aware topology may contain information on native SFs/NFs (i.e.
   SFs/NFs as known to the provider itself) and/or abstract SFs/NFs
   (i.e. logical/macro SFs/NFs representing one or more SFCs each made
   of native and/or lower level abstract SFs/NFs). As in the case of
   abstracting topology nodes, abstracting SFs/NFs is hierarchical in
   nature - the higher level of SF/NF-aware topology, the "larger"
   abstract SFs/NFs are, i.e. the larger data plane SFCs they represent.
   This allows for managing large scale networks with great number of
   SFs/NFs (such as Data Center interconnects) in a hierarchical, highly
   scalable manner resulting in control of very large number of flat in
   the data plane SFCs that span multiple domains.

11. **Dynamic Compute/VM/Storage Resource Assignment**

   In a distributed data center networks, virtual machines for compute
   resources may need to be dynamically re-allocated due to various
   reasons such as DCI network failure, compute resource load balancing,
   etc. In many cases, the DCI connectivity for the source and the
   destination is not predetermined. There may be a pool of sources and
   a pool of destination data centers associated with re-allocation of
   compute/VM/storage resources. There is no good mechanism to date to
   capture this dynamicity nature of compute/VM/storage resource
   reallocation. Generic Compute/VM/Storage resources can be described
   and announced as a SF where a DC hosting these resources can be
   modeled as an abstract node. Topology interconnecting these abstract
   nodes (DCs) in general is of multi-domain nature. Thus, SF-aware
   topology model can facilitate a joint optimization of TE network
   resources and Compute/VM/Storage resources and solve Compute/VM/
   Storage mobility problem within and between DCs.

12. **IANA Considerations**

   This document has no actions for IANA.

13. **Security Considerations**

   This document does not define networking protocols and data, hence
   are not directly responsible for security risks.

14. **Acknowledgements**

   The authors would like to thank Maarten Vissers, Joel Halpern, and
   Greg Mirsky for their helpful comments and valuable contributions.

15. **References**

15.1. **Normative References**

   [RFC7498] Quinn, P., Ed. and T. Nadeau, Ed., "Problem Statement for
              Service Function Chaining", RFC 7498,
              DOI 10.17487/RFC7498, April 2015,
              <http://www.rfc-editor.org/info/rfc7498>.

   [RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function
              Chaining (SFC) Architecture", RFC 7665,
              DOI 10.17487/RFC7665, October 2015,
              <http://www.rfc-editor.org/info/rfc7665>.

   [I-D.ietf-teas-yang-te-topo]
             Liu, X., Bryskin, I., Beeram, V., Saad, T., Shah, H., and
             O. Dios, "YANG Data Model for TE Topologies", draft-ietf-
             teas-yang-te-topo-09 (work in progress), June 2017.

   [I-D.ietf-teas-yang-te]
             Saad, T., Gandhi, R., Liu, X., Beeram, V., Shah, H., and
             I. Bryskin, "A YANG Data Model for Traffic Engineering
             Tunnels and Interfaces", draft-ietf-teas-yang-te-06 (work
             in progress), March 2017.

15.2. Informative References

   [I-D.ietf-sfc-hierarchical]
             Dolson, D., Homma, S., Lopez, D., Boucadair, M., Liu, D.,
             Ao, T., and V. Vu, "Hierarchical Service Function Chaining
             (hSFC)", draft-ietf-sfc-hierarchical-02 (work in
             progress), January 2017.

Authors' Addresses

   Igor Bryskin
   Huawei Technologies

   EMail: Igor.Bryskin@huawei.com

   Xufeng Liu
   Jabil

   EMail: Xufeng_Liu@jabil.com

   Jim Guichard
   Huawei Technologies

   EMail: james.n.guichard@huawei.com

   Young Lee
   Huawei Technologies

   EMail: leeyoung@huawei.com