

SFC WG
Internet-Draft
Intended status: Informational
Expires: January 3, 2019

D. Lachos
Unicamp
Q. Xiang
Tongji/Yale University
C. Rothenberg
Unicamp
Y. Yang
Tongji/Yale University
July 2, 2018

Multi-domain Service Function Chaining with ALTO
draft-lachos-multi-domain-sfc-alto-00

Abstract

Currently, Service Function Chaining (SFC) that span domains with different technology, administration or ownership are being defined by the SFC WG. This document focuses on how the Application Layer Traffic Optimization (ALTO) protocol can be used to advertise and discover abstract topology, resource and service information from different domains, and then compute inter-domain service function paths. Another important concern of this draft is to initiate a discussion (ALTO, SFC as well as other WGs) regarding if, how, and under what conditions ALTO can be useful to improve the multi-domain SFC process.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [RFC2119].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 3, 2019.

Copyright Notice

Copyright (c) 2018 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Context and Motivation	3
3.1. Standardization Activities	4
3.2. Research projects	4
4. ALTO for Multi-domain SFC	5
4.1. Advantages of using ALTO	5
4.1.1. Inter-domain info discovery with ALTO Property Map	6
4.1.2. Inter-domain path computation with ALTO Cost Map	6
5. Motivating Use Cases	7
5.1. ALTO as part of the SFC eXchange Platform	7
5.2. Resource Orchestration for Multi-Domain, Geo-Distributed Data Analytics	8
6. IANA Considerations	10
7. Security Considerations	10
8. Summary and Outlook	11
9. Acknowledgments	11
10. References	12
10.1. Normative References	12
10.2. Informative References	12
Authors' Addresses	14

1. Introduction

The delivery of end-to-end services often requires various Service functions (SFs) or Virtual Network Functions (VNFs). Service Function Chaining (SFC) is constructed as an abstract sequence of SFs or VNFs [[RFC7665](#)]. Multi-domain SFC is the ability to deploy SFC across multiple domains with different technology and/or

administration. To do so, an inter-domain communication process between different organizations is necessary to (i) exchange topology, resource and service information, and then (ii) compute inter-domain service function paths.

The ALTO WG has begun started to discuss the uses of ALTO as an information model for representing network resource and services in multi-domain scenarios:

- o [DRAFT-ALTO-BROKER-MDO] proposes an ALTO-based Broker-assisted architecture where a broker plane works as a coordinator between a set of top-level control planes, i.e., Multi-domain Orchestrator (MdOs). The ALTO services (with the proposed extensions) provides abstract maps with a simplified, yet enough information view about MdOs involved in the federation. This information includes the abstract network topology, resource availability (e.g., CPUs, Memory, and Storage) and capabilities (e.g., supported network functions).
- o The document [DRAFT-ALTO-UNICORN] presents Unicorn, a resource orchestration framework for multi-domain, geo-distributed data analytics. This work resorts in ALTO as the information model to support the accurate, yet privacy-preserving resource discovery across different domains. The key information to be provided by the use of ALTO including different types of resources, e.g., the computing, storage, and networking resources.

This draft offers concrete use case examples of how ALTO can be incorporated in the multi-domain SFC architecture. The examples used in this document are based on architectures and assumptions currently being discussed in the SFC WG [DRAFT-HH-MDSFC] and in the ALTO WG [DRAFT-ALTO-BROKER-MDO] [DRAFT-ALTO-UNICORN].

The overall rationale of this document is to begin a discussion between the SFC and the ALTO WG (other WGs are welcome) concerning if, how, and under which conditions ALTO will be helpful in the SFC traversing different administrative domains.

2. Terminology

This document makes use of the terminology defined in [DRAFT-HH-MDSFC], [DRAFT-ALTO-UNICORN], and [DRAFT-ALTO-BROKER-MDO].

3. Context and Motivation

Nowadays, different standardization activities (e.g., IETF and ETSI) and research projects (e.g., 5GEx [H2020.5GEX], 5G-Transformer

[[H2020-5G-TRANSFORMER](#)], T-NOVA [[T-NOVA](#)], etc.) have been focused on multi-domain network service chaining.

3.1. Standardization Activities

SFC that span domains owned by multiple administrative entities are being discussed in the IETF SFC WG. [[DRAFT-HH-MDSFC](#)], for example, describes SFC crossing different domains owned by various organizations (e.g., ISPs) or by a single organization with administration partitions. The proposed architecture uses a SFC eXchange Platform (SXP) to collect and exchange information (topology, service states, policies, etc.) between different organizations and it works both in centralized (Multiple SFC domains connected by a logical SXP) and distributed (SXP server as a broker) environments. Another IETF initiative is the Network Function Virtualization Research Group (NFVRG). The draft "Multi-domain Network Virtualization" [[DRAFT-MD-VIRT](#)] envisions a complete end-to-end logical network as stitching services offered by multiple domains from multiple providers. It also points to the need for creating solutions that enable the exchange of relevant information (resources and topologies) across different providers.

The ETSI NFV ISG is paving the way toward viable architectural options supporting the efficient placement of functions in different administrative domains. More specifically, the document [[ETSI-NFV-IFA028](#)] reports different NFV MANO architectural approaches with use cases related to network services provided using multiple administrative domains. Besides, it gives a non-exhaustive list of key information to be exchanged between administrative domains (monitoring parameters, topology view, resource capabilities, etc.) and recommendations related to security to permit the correct and proper operation of the final service.

3.2. Research projects

Several projects include an architectural model integrating NFV management with SDN control capabilities to address the challenges towards flexible, dynamic, cost-effective, and on-demand service chaining.

[[H2020.5GEX](#)] aims to integrate multiple administrations and technologies through the collaboration between operators in the context of emerging 5G networking. [[VITAL](#)][[T-NOVA](#)] follow a centralized approach where each domain advertises its capabilities to a federation layer which will act as a broker. In order to avoid one network operator per country or regions, [[H2020-5G-NORMA](#)] proposes the use of management and control into a single virtual domain. Also, the 5G-Transformer project [[H2020-5G-TRANSFORMER](#)] is defining

flexible slicing and federation of transport networking and computing resources across multiple domains.

4. ALTO for Multi-domain SFC

A "dialogue" between potential domains that will provide multi-domain SFC could be beneficial for more efficient use of resources and increasing the SFC performance. However, constrained knowledge of the network services and underlying network topology based only on localized views from the point of view of a single domain limits the potential and scope for multi-domain SFC.

To enable a highly customized multi-domains SFC, [DRAFT-HH-MDSFC] proposes a SFC eXchange Platform to realize inter-domain communication between top-level control planes. The SXP is a logical entity deployed in future Software-defined IXP (as a trusted third-party platform) or built by a single owner between different networks.

On a high level, the scope of the SXP contains two main tasks:

- o Provide end-to-end visibility through the collection of topology information, service states, and policies from different domains.
- o Compute inter-domain service function path to select the service function location from multiple candidate domains.

The ALTO protocol [RFC7285] provides abstract network information in the form of map services that can be consumed by applications in order to become network-aware and to take optimized decisions regarding traffic flows. Recently, ALTO is also being considered in multi-domain orchestration scenarios [DRAFT-ALTO-UNICORN] [DRAFT-ALTO-BROKER-MDO], in which an ALTO server can convey inter-domain network resource and topology information.

In this context, the SXP can take advantage of multi-domain ALTO services to obtain important inter-domain information to "guide" the resource/service provider selection process in that the "best" domain or candidate domains (according to established policies) can be intelligently selected. The following ALTO services can be identified:

4.1. Advantages of using ALTO

ALTO (and customized ALTO extensions) can be used to offer aggregated/abstracted views on various types of information including domain-level topology, storage resources, computation resources, networking resources and PNF/VNF capabilities. This generic

representation contributing to a more simple and scalable solution for resource and service discovery in multi-domain, multi-technology environments.

In case of Multi-domain SFC, the following ALTO services could be identified:

4.1.1. Inter-domain info discovery with ALTO Property Map

Each domain needs a global view of other potential candidate domains to know who can provide part of the NF in the SFC. A brief list of information to be exchanged between different domains includes:

- o Resource capabilities, applicable to both IT (computing and storage) and networking resources participant of the multi-domain SFC, to assist on the decision of SFs placement.
- o Access information (e.g., URL) to the orchestrator entry points and Service Access Points (SAPs) for a corresponding network/domain.

The ALTO Property Map Service [[DRAFT-ALTO-PM](#)] can provide a clear global view of the resource information offered by other domains. This information allows discovering which candidate domains may be contacted to deliver the remaining requirements of a requested end-to-end service deployment.

4.1.2. Inter-domain path computation with ALTO Cost Map

Once the candidate domains are discovered, it is necessary to compute inter-domain service function path to select the service function location from those different candidate domains.

The connectivity information among discovered domains can be retrieved by an ALTO Cost Map service, responding, for instance, a path vector with the AS-level topology distance between the source domain and candidate domains. Moreover, path vector constraints (as described in the Multi-Cost Map [[RFC8189](#)]) can be applied to filter out the list of unqualified domains.

In case of the Hybrid Hierarchical SFC architecture [[DRAFT-HH-MDSFC](#)], the SXP (or the Path Calculation Element in the top-level control plane) could use this information to compute multi-domain service function paths.

5. Motivating Use Cases

5.1. ALTO as part of the SFC eXchange Platform

As mentioned earlier, the draft [DRAFT-HH-MDSFC] defines a multi-domain SFC architecture that combines control planes to be deployed either into a large domain consisting of smaller sub-domains owned by the same organization or into multiple large domains with different ownership. Figure 1 shows a SXP connecting three different domains (AS1, AS2, AS3). Each domain provides different SFs: AS1 -> SF1; AS2 -> SF2 and SF3; AS3 -> SF3. The SXP includes an ALTO server component to provide abstract topology, resource, and service information for the high-level control plane in each domain.

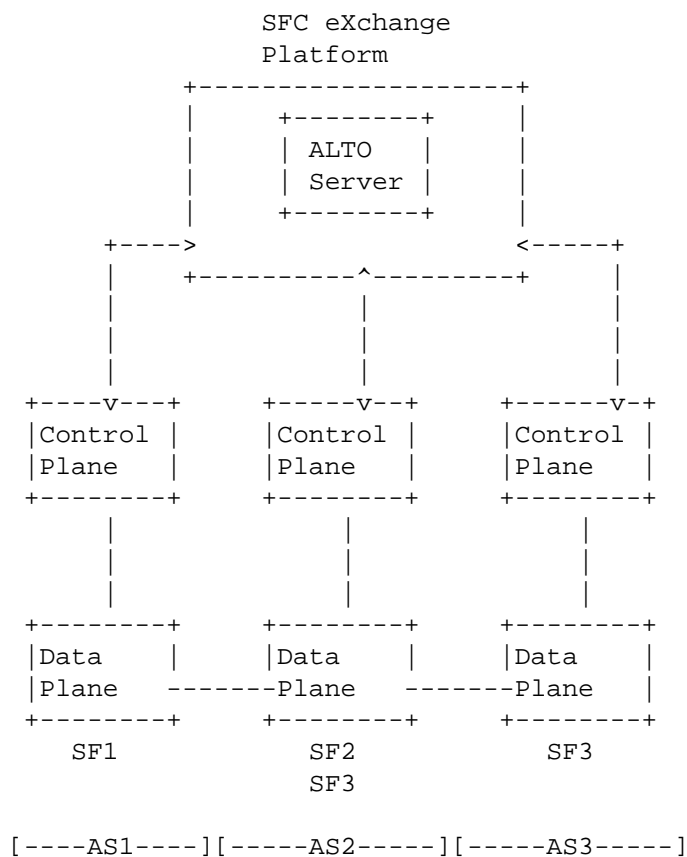


Figure 1: ALTO as part of the SFC eXchange Platform

Every domain has a local Information Base Element; this component can be used by the SXP to create hierarchical databases containing inter-domain resource and topology information. This information source is

used by the ALTO server to create two different ALTO Map Services: (i) Property Map and (ii) Cost Map.

The Property Map includes a property value grouped by Autonomous System (AS), this value contains the supported network functions. Additional properties could be considered such as resource availability (e.g., CPUs, Memory, and Storage), orchestrator entry points, etc. An example of the Property Map in our basic scenario is:

	Capabilities	Entry Point	CPU	MEM	Storage	...
AS1	{SF1}	http://...
AS2	{SF2, SF3}	http://...
AS3	{SF3}	http://...

Table 1: ALTO Property Map

The Cost Map defines a path vector as an array of ASes, representing the AS-level topological distance for a given SFC request. Table 2 below shows a brief example of a service request and its inter-domain service function path response containing a list of potential domains to be traversed to deliver such service.

SFC Request	Multi-domain Service Function Path(s)
SF1->SF2->SF3	1: {AS1:SF1->AS2:SF2->AS2:SF3}
	2: {AS1:SF1->AS2:SF2->AS3:SF3}

Table 2: ALTO Cost Map

5.2. Resource Orchestration for Multi-Domain, Geo-Distributed Data Analytics

In addition to commercial SFC, ALTO is also used as a core information model for collaborative data science networks. The document [DRAFT-ALTO-UNICORN] presents the design of Unicorn, a unified resource orchestration framework for multi-domain, geo-distributed data analytics, currently being developed and deployed in the CMS network, one of the largest scientific experiments in the LHC network.

ALTO is well suited as a fundamental component in Unicorn for providing a generic representation that (1) allows different types of

data analytics jobs to accurately describe their resource requirements and (2) allows member networks to provide accurate information on different types of resources they own and at the same time maintain their privacies.

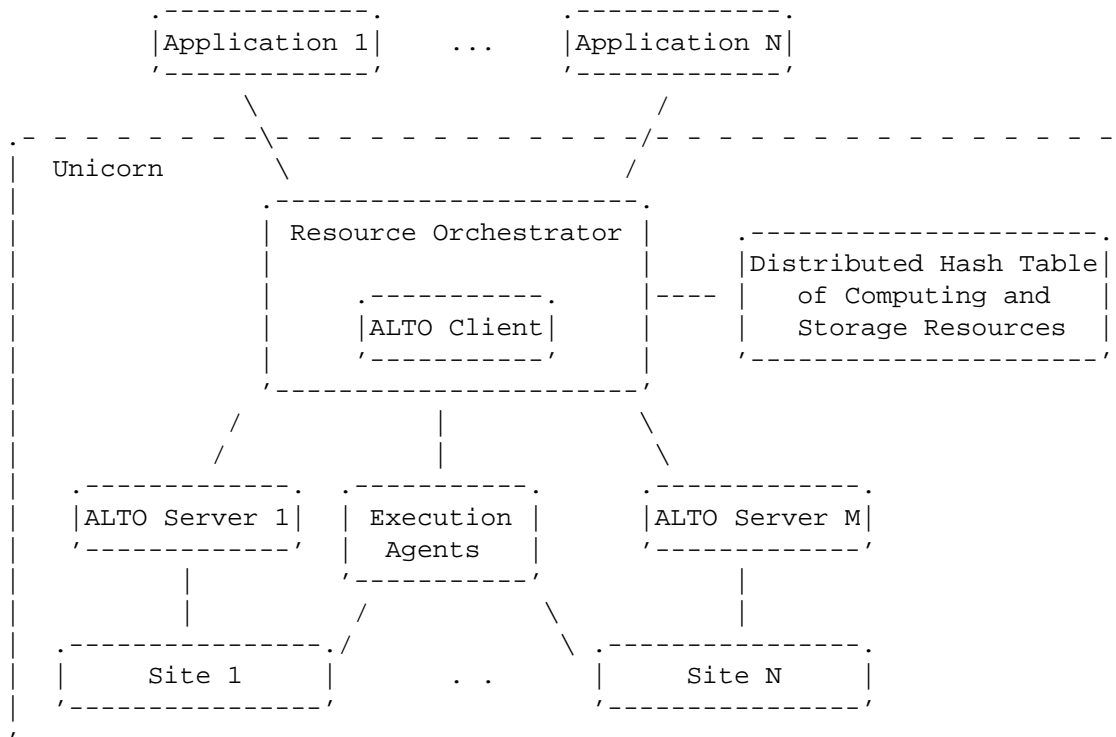


Figure 2: Architecture of Unicorn.

Figure 2 presents the architecture of Unicorn. Specifically, for each member network, one or more ALTO servers are deployed to provide accurate, yet privacy-preserving information of different types of resources owned by the corresponding network. Examples of such information include the link bandwidth between endpoints, the memory I/O bandwidth and the CPU utilization at computing endpoints and the storage space at storage endpoints. In addition to the basic ALTO services defined in [RFC7285], The ALTO servers in Unicorn also provide ALTO extension services such as the ALTO Multi-Cost Service [RFC8189], the ALTO Server-Sent Event Service [DRAFT-ALTO-INCR-UPD] and the ALTO Path Vector Service [DRAFT-ALTO-PV] to provide fine-grained resource information.

Because the ALTO Path Vector service may expose additional private information of each network, Unicorn develops an obfuscating protocol

which ensures that nor the orchestrator or any member networks can associate any path vector information with a corresponding network.

To better address the scalability issue of multi-domain resource discovery, Unicorn also develops a proactive full-mesh discovery mechanism, which precomputes network-level ALTO path vector information and performs projection using such information to compute the fine-grained resource information in response to orchestrator's resource discovery requests.

Details of the obfuscating protocol and the proactive full-mesh discovery mechanism developed in Unicorn can be found in the [\[DRAFT-ALTO-UNICORN\]](#) document.

6. IANA Considerations

This document includes no request to IANA.

7. Security Considerations

The ALTO base protocol has an extensive discussion on potential security and privacy issues. Using the ALTO base protocol to support multi-domain SFC will not raise new security and privacy issue. However, the information provided by the ALTO base protocol is considered coarse-grained in several recent use cases. As a result, several ALTO extension services have been designed to provide fine-grained network information to the application. Using these ALTO extension services for multi-domain SFC would raise new security and privacy concerns. Next, we list these issues on a per extension basis.

The ALTO unified property extension [\[DRAFT-ALTO-PM\]](#) generalizes the concept of endpoint properties to other entity domains, such as abstract network element. The properties of these entities may contain sensitive service-function-specific information. Exposing such information may discourage networks to provide fine-grained information to support multi-domain SFC.

The ALTO performance cost metrics extension [\[DRAFT-ALTO-METRICS\]](#) proposes a set of ALTO cost metrics derived from traffic engineering tools and protocols. It is stated in this extension that "sharing network TE metric values in numerical mode requires full mutual confidence between the entities managing the ALTO Server and Client." In multi-domain SFC use case, such mutual confidence is needed not only between ALTO server and client, but also among all networks, and third-parties such as broker and a global orchestrator. How to achieve such mutual confidence in multi-domain SFC use case requires further investigation.

The ALTO path vector extension [DRAFT-ALTO-PV] allows ALTO clients to query network information such as capacity region for a given set of flows. Several related studies have shared concerns that this extension may reveal more network internal structures than the more abstract single-node abstraction used in the ALTO base protocol. In multi-domain SFC, this concern will further be amplified as third-party participants may access such information. The recent designed Unicorn system proposes an obfuscating protocol that prevents the receiver of the capacity region information from associating this region to any network. This protocol sheds light for addressing the privacy issue brought by the ALTO path vector extension.

The ALTO cost calendar [DRAFT-ALTO-CALENDAR] and the ALTO incremental update [DRAFT-ALTO-INCR-UPD] extensions allow the ALTO client to get temporal network information. The intention of these extensions is to allow applications to make flexible decisions on when to use network information. However, both extensions expose temporal policy and traffic information of network so that a user may know when the network is most vulnerable for overloading. This issue needs to be carefully addressed in order for both extensions to be used for multi-domain SFC.

8. Summary and Outlook

This draft provided arguments why ALTO is a meaningful protocol in the context of SFC traversing different domains, and it presented use case examples about the how ALTO can be used to advertise and discover abstract topology, resource and service information from different domains, and then compute inter-domain service function paths.

The overall objective of this document is to arouse discussions in the SFC WG in order to assess the suitability of the ALTO as a useful protocol for multi-domain SFC scenarios. The result of such discussions will be captured in future versions of this draft.

9. Acknowledgments

This work is supported by the Innovation Center of Ericsson S.A., Brazil (grant agreement UNI.64).

Many thanks to Sabine Randriamasy, and Lyle Bertz for their feedback on this draft.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC7285] Alimi, R., Ed., Penno, R., Ed., Yang, Y., Ed., Kiesel, S., Previdi, S., Roome, W., Shalunov, S., and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol", [RFC 7285](#), DOI 10.17487/RFC7285, September 2014, <<https://www.rfc-editor.org/info/rfc7285>>.
- [RFC7665] Halpern, J., Ed. and C. Pignataro, Ed., "Service Function Chaining (SFC) Architecture", [RFC 7665](#), DOI 10.17487/RFC7665, October 2015, <<https://www.rfc-editor.org/info/rfc7665>>.
- [RFC8189] Randriamasy, S., Roome, W., and N. Schwan, "Multi-Cost Application-Layer Traffic Optimization (ALTO)", [RFC 8189](#), DOI 10.17487/RFC8189, October 2017, <<https://www.rfc-editor.org/info/rfc8189>>.

10.2. Informative References

- [DRAFT-ALTO-BROKER-MDO]
Perez, D. and C. Rothenberg, "ALTO-based Broker-assisted Multi-domain Orchestration", [draft-lachosrothenberg-alto-brokermdo-00](#) (work in progress), March 2018.
- [DRAFT-ALTO-CALENDAR]
Randriamasy, S., Yang, Y., Wu, Q., Lingli, D., and N. Schwan, "ALTO Cost Calendar", [draft-ietf-alto-cost-calendar-05](#) (work in progress), June 2018.
- [DRAFT-ALTO-INCR-UPD]
Roome, W., Yang, Y., and S. Chen, "ALTO Incremental Updates Using Server-Sent Events (SSE)", [draft-ietf-alto-incr-update-sse-11](#) (work in progress), June 2018.
- [DRAFT-ALTO-METRICS]
Wu, Q., Yang, Y., Lee, Y., Dhody, D., and S. Randriamasy, "ALTO Performance Cost Metrics", [draft-ietf-alto-performance-metrics-04](#) (work in progress), June 2018.

[DRAFT-ALTO-PM]

Roome, W., Chen, S., Randriamasy, S., Yang, Y., and J. Zhang, "Unified Properties for the ALTO Protocol", [draft-ietf-alto-unified-props-new-03](#) (work in progress), March 2018.

[DRAFT-ALTO-PV]

Bernstein, G., Chen, S., Gao, K., Lee, Y., Roome, W., Scharf, M., Yang, Y., and J. Zhang, "ALTO Extension: Path Vector Cost Type", [draft-ietf-alto-path-vector-03](#) (work in progress), March 2018.

[DRAFT-ALTO-UNICORN]

Xiang, Q., Le, F., Yang, Y., Newman, H., and H. Du, "Unicorn: Resource Orchestration for Multi-Domain, Geo-Distributed Data Analytics", [draft-xiang-alto-multidomain-analytics-01](#) (work in progress), March 2018.

[DRAFT-HH-MDSFC]

Li, G., Li, G., Li, T., Xu, Q., and H. Zhou, "Hybrid Hierarchical Multi-Domain Service Function chaining", [draft-li-sfc-hhsfc-04](#) (work in progress), April 2018.

[DRAFT-MD-VIRT]

Bernardos, C., Contreras, L., Vaishnavi, I., Szabo, R., Li, X., Paolucci, F., Sgambelluri, A., Martini, B., Valcarengi, L., Landi, G., Andrushko, D., and A. Mourad, "Multi-domain Network Virtualization", [draft-bernardos-nfvrg-multidomain-04](#) (work in progress), March 2018.

[ETSI-NFV-IFA028]

ETSI, "Report on architecture options to support multiple administrative domains V3.1.1", Jan 2018, http://www.etsi.org/deliver/etsi_gr/NFV-IFA/001_099/028/03.01.01_60/gr_NFV-IFA028v030101p.pdf.

[H2020-5G-NORMA]

H2020, "5G-NORMA -- 5G Novel Radio Multiservice adaptive network Architecture", 2015, <https://5gnorma.5g-ppp.eu/>.

[H2020-5G-TRANSFORMER]

H2020, "5G-Transformer -- 5G Mobile Transport Platform for Vertical", 2017, <http://5g-transformer.eu/>.

[H2020.5GEX]

H2020, "5GEx -- 5G Exchange Project", 2014, <http://www.5gex.eu/>.

[T-NOVA] FP7 project T-NOVA, "T-NOVA Project, Network Functions as a Service over Virtualised Infrastructures", 2014, <<http://www.t-nova.eu/>>.

[VITAL] VITAL PROJECT H2020, "VITAL -- Virtualized hybrid satellite-Terrestrial systems for resilient and flexible future networks", 2015, <<http://www.ict-vital.eu/>>.

Authors' Addresses

Danny Alex Lachos Perez
University of Campinas
Av. Albert Einstein 400
Campinas, Sao Paulo 13083-970
Brazil

Email: dlachosp@dca.fee.unicamp.br
URI: <https://intrig.dca.fee.unicamp.br/danny-lachos/>

Qiao Xiang
Tongji/Yale University
51 Prospect Street
New Haven, CT
USA

Email: qiao.xiang@cs.yale.edu

Christian Esteve Rothenberg
University of Campinas
Av. Albert Einstein 400
Campinas, Sao Paulo 13083-970
Brazil

Email: chesteve@dca.fee.unicamp.br
URI: <https://intrig.dca.fee.unicamp.br/christian/>

Y. Richard Yang
Tongji/Yale University
51 Prospect St
New Haven, CT
USA

Email: yang.r.yang@gmail.com