

ALTO WG
Internet-Draft
Intended status: Informational
Expires: January 13, 2022

LM. Contreras
Telefonica
D. Lachos
C. Rothenberg
Unicamp
S. Randriamasy
Nokia Bell Labs
July 12, 2021

**Use of ALTO for Determining Service Edge
draft-contreras-alto-service-edge-03**

Abstract

Service providers are starting to deploy and interconnect computing capabilities across the network for hosting network functions and applications. In distributed computing environments, both computing and topological information are necessary in order to determine the more convenient infrastructure where to deploy such a service or application. This document proposes an initial approach towards the use of ALTO to provide such information and assist the selection of appropriate deployment locations for services and applications.

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 13, 2022.

Copyright Notice

Copyright (c) 2021 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.	Computing needs	3
3.	Usage of ALTO for determining where to deploy a function or application	4
3.1.	Integrating compute information in ALTO	5
3.2.	Association of compute capabilities to network topology .	5
3.3.	ALTO architecture for determining serve edge	5
4.	Definition of flavors in ALTO property map	6
5.	IANA Considerations	7
6.	Security Considerations	7
7.	Conclusions	7
8.	References	8
8.1.	Normative References	8
8.2.	Informative References	8
	Authors' Addresses	9

[1.](#) Introduction

The advent of virtualization is enabling the operators with a dynamic instantiation of network functions and applications by using different techniques on top of commoditized computation infrastructures, permitting a flexible and on-demand deployment of services, aligned with the actual needs observed as demanded by the customers.

Operators are starting to deploy distributed computing environments in different parts of the network with the objective of addressing the different service needs in terms of latency, bandwidth, processing capabilities, etc. This is translated in the emergence of a number of data centers of different sizes (e.g., large, medium, small) characterized by distinct dimension of CPUs, memory and storage capabilities, as well as bandwidth capacity for forwarding the traffic generated in and out the corresponding data center.

The probable future situation, with the generalization and proliferation of the edge computing approach, will increase the potential footprint where a function or application can be deployed. These different dimensioning rules result in a different unitary cost

per CPU, memory, and storage in each computing environment because of the scale.

All the available distributed computing capabilities can complicate the decision of what infrastructure use for instantiating a given function or application. Such a decision influences not only the resources that are consumed in a given computing environment, but also the network capacity of the path that connects such environment with the rest of the network from traffic source to destination.

It is then essential for a network operator to have mechanisms assisting on the decision by considering a number of constraints related to the function or application to be deployed understanding how a given decision on the computing environment for the service edge affects to the transport network substrate. This would allow to integrate network capabilities in the function placement decision and further optimize performance of the deployed application.

This document proposes the usage of ALTO [[RFC7285](#)] for assisting with such a decision.

2. Computing needs

A given network function or application typically shows certain requirements in terms of processing capabilities (i.e., CPU), as well as volatile memory (i.e., RAM) and storage capacity.

Cloud computing providers, such as Amazon Web Services or Microsoft Azure, typically structure their offerings of computing capabilities by bundling CPU, RAM and storage units as quotas, instances or flavors that can be consumed in an ephemeral or temporal fashion, during the actual lifetime of the required function or application.

This same approach is being taken nowadays for characterizing bundles of resources on the so-called Network Function Virtualization Infrastructure (NFVI) Points of Presence (PoPs) being deployed by the telco operators. Specifically, the Common Network Function Virtualisation Infrastructure Telecom Taskforce (CNTT) [[CNTT](#)], [[GSMA](#)], jointly hosted by GSMA and the Linux Foundation, is intending to harmonize the definition of above-mentioned computing capability instances or flavors for abstracting capabilities of the underlying NFVI facilitating a more efficient utilization of the infrastructure and simplifying the integration and certification of functions, where certification means the assessment of the expected behavior for a given function according to the level of resources determined by a given flavor.

Focusing on the CNTT ongoing work, the flavors or instances are characterized according to:

- o Type of instance (T): the types of instances are characterized as B (Basic), or N (Network Intensive). The latter can come with extensions for network acceleration for offloading network intensive operations to hardware.
- o Interface Option (I): it refers to the associated bandwidth of the network interface.
- o Compute flavor (F): it refers to a given predefined combination of resources in terms of virtual CPU, RAM, disk, and bandwidth for the management interface.
- o Optional storage extension (S): allows to request additional storage capacity.
- o Optional hardware acceleration characteristics (A): to request specific acceleration capabilities for improving the performance of the infrastructure.

The naming convention of an instance is thus encoded as T.I.F.S.A.

3. Usage of ALTO for determining where to deploy a function or application

ALTO can assist the deployment of a service or application on a specific flavor or instance of the computing substrate by taking into consideration network cost metrics.

A generic and primary approach is to take into account metrics related to the computing environment, such as availability of resources, unitary cost of those resources, etc.

Nevertheless, the function or application to be deployed on top of a given flavor is interconnected outside the computing environment where it is deployed, also requiring to guarantee transport network requirements to ensure the application performance, such as bandwidth, latency, etc.

The objective then is to leverage on ALTO to provide information about the more convenient execution environments to deploy virtualized network functions or applications, allowing the operator to get a coordinated service edge and transport network recommendation.

3.1. Integrating compute information in ALTO

CNTT proposes the existence of a catalogue of compute infrastructure profiles collecting the computing capability instances available to be consumed. Such kind of catalogue could be communicated to ALTO or even incorporated to it.

ALTO server queries are required to support T.I.F.S.A encoding in order to retrieve proper maps from ALTO. Additionally, filtered queries for particular characteristics of a flavor could also be supported.

3.2. Association of compute capabilities to network topology

It is required to associate the location of the available instances with topological information to allow ALTO construct the overall map. The expectation is to manage the network and cloud capabilities by the same entity, handling both network and compute abstractions jointly, producing an integrated map.

At this stage three potential solutions could be considered:

- o To leverage on (and possibly extend) [[I-D.ietf-teas-sf-aware-topo-model](#)] for disseminating topology information together with notion of function location (that would require to be adapted to the existence of available compute capabilities). A recent effort in this direction can be found in [[I-D.llc-teas-dc-aware-topo-model](#)].
- o To extend BGP-LS [[RFC7752](#)], which is already considered as mechanism for feeding topology information in ALTO, in order to also advertise computing capabilities as well.
- o To combine information from the infrastructure profiles catalogue with topological information by leveraging on the IP prefixes allocated to the gateway providing connectivity to the NFVI PoP.

The viability of these options will be explored in future versions of this document.

3.3. ALTO architecture for determining serve edge

The following logical architecture defines the usage of ALTO for determining service edges.

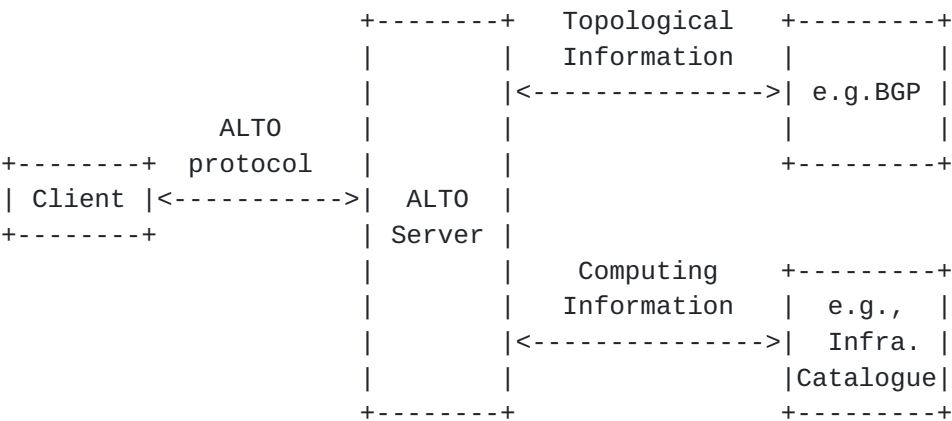


Figure 1: Service Edge Information Exchange

4. Definition of flavors in ALTO property map

The ALTO unified property extension [DRAFT-PM] generalizes the concept of endpoint properties to domains of other entities through property maps. In the context of the CNTT domain, an ALTO property map could be used to expose T.I.F.S.A information (.e., potential candidate flavors) of available NFVI PoPs where an application or service can be deployed.

Table 1 below shows an illustrative example of an ALTO property map with property values grouped by flavor name.

Flavor Name	Type of instance (T)	Interface Option (I)	Compute flavor (F) {CPU, RAM, disk and bandwidth}	S.	A.
Small-1	Basic	{1, 2, 3, 4, 5, 6, 7, 8, 9 Gbps}	{1,512 MB,1 GB,1 Gbps}
Small-2	Network Intensive	{1, 2, 3, 4, 5, 6, 7, 8, 9 Gbps}	{1,512 MB,1 GB,1 Gbps}
Medium-1	Network Intensive	{25, 50, 75, 100, 125, 150 Gbps}	{2,4 GB,40 GB,1 Gbps}
Large-1	Compute Intensive	{50, 100, 150, 200, 250, 300 Gbps}	{4,8 GB,80 GB,1 Gbps}
Large-2	Compute Intensive	{100, 200, 300, 400, 500, 600 Gbps}	{8,16 GB,160 GB,1 Gbps}
...

Table 1: ALTO Property Map

5. IANA Considerations

This document includes no request to IANA.

6. Security Considerations

TBD.

7. Conclusions

Telco networks will increasingly contain a number of interconnected data centers, of different size and characteristics, allowing flexibility in the dynamic deployment of functions and applications for advance services. The overall objective of this document is to begin a discussion in the ALTO WG regarding the suitability of the ALTO protocol for determining where to deploy a function or application in distributed computing environments. The result of such discussions will be reflected in future versions of this draft.

8. References

8.1. Normative References

- [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>>.
- [RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", [RFC 7752](#), DOI 10.17487/RFC7752, March 2016, <<https://www.rfc-editor.org/info/rfc7752>>.
- [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>>.

8.2. Informative References

- [CNTT] "Cloud infrastructure Telco Taskforce Reference Model, Version 6.0", January 2021, <https://cntt-n.github.io/CNTT/doc/ref_model/>.
- [DRAFT-PM] Roome, W., Randriamasy, S., Yang, Y., Zhang, J., and K. Gao, "Unified Properties for the ALTO Protocol", [draft-ietf-alto-unified-props-new-17](#) (work in progress), April 2021.
- [GSMA] "Cloud Infrastructure Reference Model, Version 1.0", November 2020, <<https://www.gsma.com/newsroom/wp-content/uploads//NG.126-v1.0-2.pdf>>.
- [I-D.ietf-teas-sf-aware-topo-model] Bryskin, I., Liu, X., Lee, Y., Guichard, J., Contreras, L., Ceccarelli, D., and J. Tantsura, "SF Aware TE Topology YANG Model", [draft-ietf-teas-sf-aware-topo-model-08](#) (work in progress), July 2021.
- [I-D.llc-teas-dc-aware-topo-model] Lee, Y., Liu, X., and L. Contreras, "DC aware TE topology model", [draft-llc-teas-dc-aware-topo-model-00](#) (work in progress), November 2020.

Authors' Addresses

Luis M. Contreras
Telefonica
Ronda de la Comunicacion, s/n
Madrid 28050
Spain

Email: luismiguel.contrerasmurillo@telefonica.com

URI: <http://lmcontreras.com/>

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/>

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/>

Sabine Randriamasy
Nokia Bell Labs
Route de Villejust
Nozay 91460
France

Email: sabine.randriamasy@nokia-bell-labs.com

