

NMRG
Internet-Draft
Intended status: Informational
Expires: 27 April 2023

LM. Contreras
Telefonica
P. Lucente
NTT
October 2022

Interconnection Intents
draft-contreras-nmrg-interconnection-intents-03

Abstract

This memo introduces the use case of the usage of intents for expressing advance interconnection features, further than traditional IP peering.

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 4 April 2023.

Copyright Notice

Copyright (c) 2022 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 Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction	2
2.	Evolution of Network interconnection	3
2.1.	Potential interconnection intent types	3
2.2.	Interconnection intent lifecycle	4
2.3.	Protocol aspects	5
3.	Interconnection intent structure	5
4.	Security Considerations	6
5.	IANA Considerations	6
6.	References	6
	Acknowledgments	7
	Authors' Addresses	7

[1.](#) Introduction

The success of Internet-based services has been built on top of the global reachability of content accessed by the end-users, which is facilitated by the interconnection of individual networks owned by distinct service providers constituting independent administrative domains.

Such interconnection services have been initially based simply on delivery of IP traffic between the interconnected parties leveraging on BGP. This peer model enables full connectivity. However, the traditional interconnection model shows some limitations when additional information to that related to routing is needed.

New network capabilities based on programmability and virtualization are producing service situations where a connectivity-only approach is not sufficient. The increasing availability of computing capabilities internal to the networks, or attached to them, enable new scenarios where those capabilities can be consumed through the advertisement or exposure of these execution environments (i.e., in terms of compute, storage and associated networking resources). Such information from an interconnected provider can be obtained from e.g. [[I-D.llc-teas-dc-aware-topo-model](#)].

In addition or complementary to that, even services or network functions could be advertised in order to make them available for interconnection. For example, as service we could consider the advertisement of CDN capabilities as in CDNi approach [[RFC7336](#)], while as network function we could consider functions like firewall, CGNAT, etc, present in the network [[I-D.ietf-teas-sf-aware-topo-model](#)].

All these scenarios present clear evolutions of the interconnection model which can not be simply expressed through existing mechanisms, or at least, cannot be expressed in a simple (and comprehensive) way with such existing mechanisms. Here is where an advanced interconnection intent can assist on declaring the goal of the interconnection transcending pure IP traffic exchange and including more advance capabilities as the ones mentioned before.

2. Evolution of Network interconnection

It becomes clear the trend to increasingly rely on multi-domain scenarios for the provision of services. For instance, the access today to an on-demand OTT video on Internet implies the interaction of more than one single administrative domain. Thus, end-to-end service delivery over multiple providers or domains is becoming the norm.

Complex network services leveraging on virtualization solutions and different infrastructure environments pertaining to distinct administrative domains (i.e., operated and managed by distinct providers) can be easily foreseen.

It is then necessary to explore mechanisms for interconnecting that multiple domain environments in a common, portable way independently of the owner of such infrastructure.

2.1. Potential interconnection intent types

The interconnection intent should provide enough abstractions to express a variety of interconnection options.

The purpose of the interconnection intent can be multiple:

- * To enable multi-domain network service programming, by soliciting interconnection of service / network functions in different domains
- * To enable multi-domain deployment of virtualized network functions, by advertising the availability of compute and storage resources in different domains
- * To facilitate multi-domain network function or service charging, by advertising (cumulative) costs in the different domains
- * To enable traffic interchange, ie. IP as in traditional peering or optical

- * To put in place the right collection of policies to implement and operate the interconnection
- * To facilitate whatever combination of all of them

2.2. Interconnection intent lifecycle

[RFC9315] defines an intent lifecycle composed of two phases, namely fulfillment and assurance. Figure 1 captures the intent procedure for the fulfillment phase.

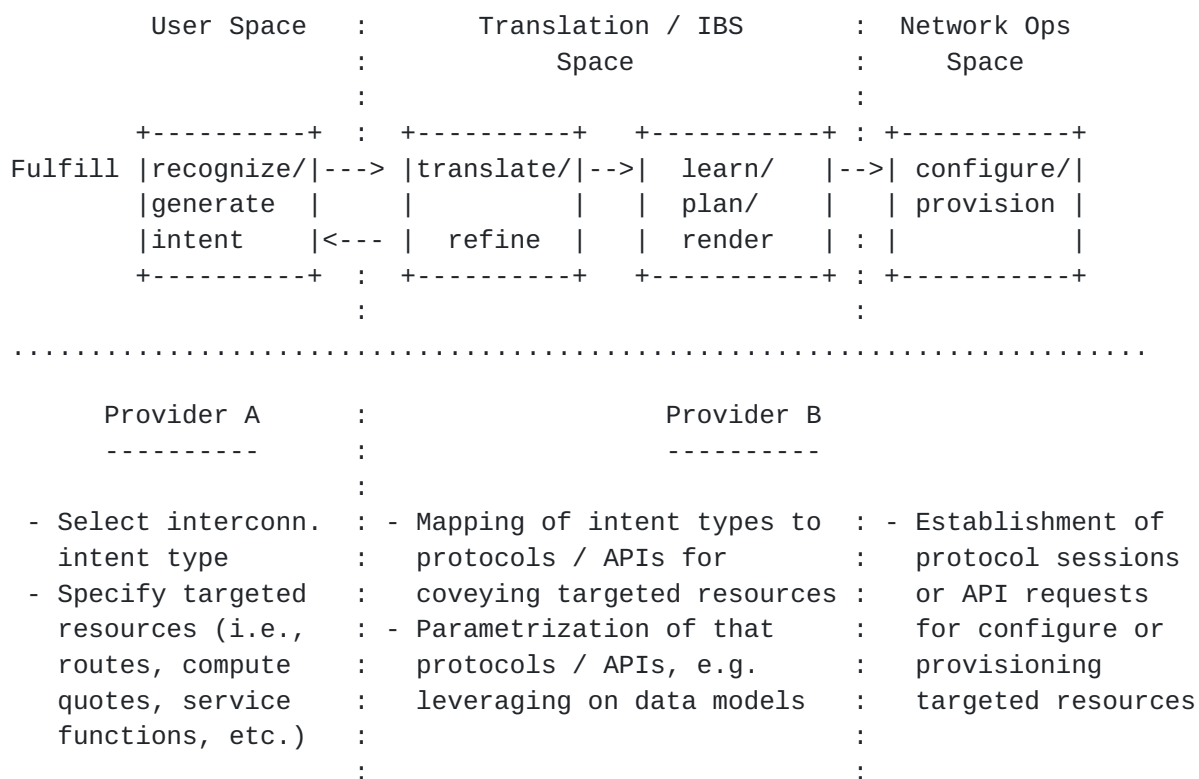


Figure 1: Fulfillment phase of the Interconnection Intent

Similarly, Figure 2 sketches the intent procedure for the assurance phase.

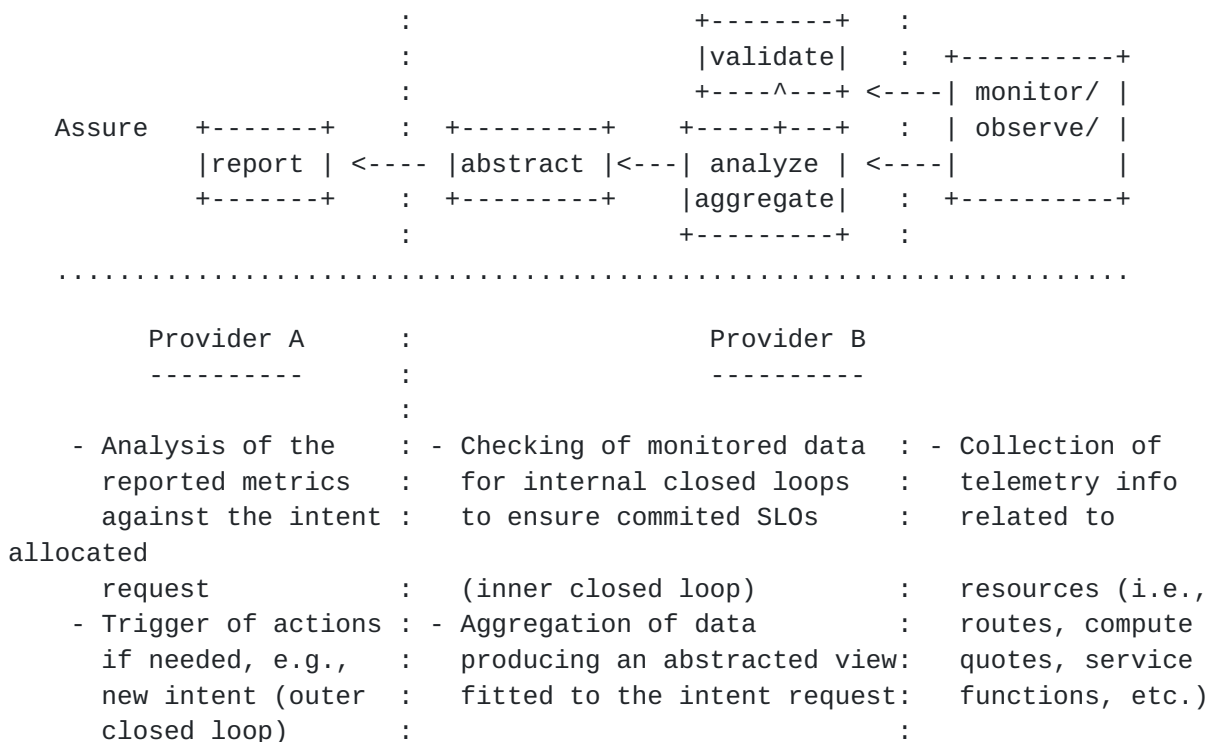


Figure 2: Assurance phase of the Interconnection Intent

Both Fulfillment and Assurance phases are integral part of the interconnection intent.

2.3. Protocol aspects

Ultimately the ideas and notions elaborated in this document will need to find room in a framework made of one or multiple protocols (ie. BGP, LISP, ALTO, etc.) and/or API definitions. While the exact definition of such framework is left as future work, in this document we intend to perform some seminal work in this sense (ie. identify existing protocols that could fit, determine gaps of such protocols, etc.).

3. Interconnection intent structure

In order to address the different interconnection intent types described in [section 2.1](#), the structure of the intent should be sufficiently flexible to allow the expression of different targets. Thus, the intent structure could include:

- * Information of the type of data traffic being subject of the interconnection intent (e.g., IP prefixes involved) among providers.

- * Service functions expected to be supported by the peer provider. These could be expressed in terms of type of service function and number of instances required. Furthermore, it can be necessary to consider how the service functions are expected to be connected in terms of topology (i.e., service function graph).
- * Resources expected to be offered by the peer provider. These could be expressed in terms of raw values of number of CPUs, memory and storage size, or bandwidth capacity, or alternatively, in terms of quotas grouping resources in a predefined manner.
- * Constraints that could apply to whatever of the elements included in the interconnection intent, including traffic steering ones. Aspects such as committed rates, burst size, cumulative traffic, service function affinity, redundancy, traffic engineering (e.g., latency), etc., could be part of such constraints.
- * Further information that could be necessary for delivering an end-to-end service by means of the intent.

4. Security Considerations

To be done.

5. IANA Considerations

This draft does not include any IANA considerations

6. References

[I-D.ietf-teas-sf-aware-topo-model]

Bryskin, I., Liu, X., Lee, Y., Guichard, J., Contreras, L. M., Ceccarelli, D., Tantsura, J., and D. Shytyi, "SF Aware TE Topology YANG Model", Work in Progress, Internet-Draft, [draft-ietf-teas-sf-aware-topo-model-09](https://www.ietf.org/archive/id/draft-ietf-teas-sf-aware-topo-model-09), 27 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-teas-sf-aware-topo-model-09.txt>>.

[I-D.llc-teas-dc-aware-topo-model]

Lee, Y., Liu, X., and M. Luis Contreras, "DC aware TE topology model", Work in Progress, Internet-Draft, [draft-llc-teas-dc-aware-topo-model-02](https://www.ietf.org/archive/id/draft-llc-teas-dc-aware-topo-model-02), 11 July 2022, <<https://www.ietf.org/archive/id/draft-llc-teas-dc-aware-topo-model-02.txt>>.

- [RFC7336] Peterson, L., Davie, B., and R. van Brandenburg, Ed., "Framework for Content Distribution Network Interconnection (CDNI)", [RFC 7336](#), DOI 10.17487/RFC7336, August 2014, <<https://www.rfc-editor.org/info/rfc7336>>.
- [RFC9315] Clemm, A., Ciavaglia, L., Granville, L. Z., and J. Tantsura, "Intent-Based Networking - Concepts and Definitions", [RFC 9315](#), DOI 10.17487/RFC9315, October 2022, <<https://www.rfc-editor.org/info/rfc9315>>.

Acknowledgments

This work has been partly funded by the European Commission through the H2020 project 5GROWTH (Grant Agreement no. 856709).

Authors' Addresses

Luis M. Contreras
Telefonica
Ronda de la Comunicacion, s/n
Sur-3 building, 1st floor
28050 Madrid
Spain
Email: luismiguel.contrerasmurillo@telefonica.com
URI: <http://lmcontreras.com/>

Paolo Lucente
NTT
Siriusdreef 70-72
2132 Hoofddorp, WT
Netherlands
Email: paolo@ntt.net

