



Exposure of Telefonica network topology through ALTO for integration with Telefonica CDN Update from IETF 114

Luis M. Contreras (*) | *Telefónica GCTIO – Transport Group*
Francisco Cano, Anais Escribano | *Telefónica CCDO – Video Group*

MOPS WG meeting @ IETF 115, London, November 2022

Rationale for making TCDN to be transport network aware (reminder)

- One of the main objective of TCDN is to provide an efficient delivery of contents within the network
- Content delivery is based nowadays on a (semi-)static view of the network, decoupled from the real situation along time
- In order to make a complete and efficient usage of the network, TCDN would benefit from a real time knowledge of the status and characteristics of the network
 - For instance, allowing delivery decisions in TCDN to quickly adapt to network status variation (e.g., topology changes, congestion, etc.)
- Project presented in IETF 114, update reported here

Network map & Cost map (reminder)

Networkmap – association of prefixes per PID ↔ BGP
Costmap - hopcount among PIDs ↔ BGP-LS

Network Map

```
{
  "pid0:0a0a0a01": [
    "1.1.1.0/24"
  ],
  "pid0:0a0a0a02": [
    "2.2.2.0/24"
  ],
  "pid0:0a0a0a03": [
    "3.3.3.0/24"
  ],
  "pid0:0a0a0a05": [
    "11.11.11.0/30"
  ],
  "pid0:0a0a0a06": [
    "22.22.22.0/30"
  ]
}
```

Cost Map

```
{
  "pid0:0a0a0a01": {
    "pid0:0a0a0a01": 0,
    "pid0:0a0a0a02": 2,
    "pid0:0a0a0a03": 2,
    "pid0:0a0a0a05": 2,
    "pid0:0a0a0a06": 1
  },
  "pid0:0a0a0a02": {
    "pid0:0a0a0a01": 2,
    "pid0:0a0a0a02": 0,
    "pid0:0a0a0a03": 2,
    "pid0:0a0a0a05": 2,
    "pid0:0a0a0a06": 1
  },
  "pid0:0a0a0a03": {
    "pid0:0a0a0a01": 2,
    "pid0:0a0a0a02": 2,
    "pid0:0a0a0a03": 0,
    "pid0:0a0a0a05": 2,
    "pid0:0a0a0a06": 1
  },
  "pid0:0a0a0a05": {
    "pid0:0a0a0a01": 2,
    "pid0:0a0a0a02": 2,
    "pid0:0a0a0a03": 2,
    "pid0:0a0a0a05": 0,
    "pid0:0a0a0a06": 3
  },
  "pid0:0a0a0a06": {
    "pid0:0a0a0a01": 1,
    "pid0:0a0a0a02": 1,
    "pid0:0a0a0a03": 1,
    "pid0:0a0a0a05": 3,
    "pid0:0a0a0a06": 0
  }
}
```

A PID can represent a set of prefixes assigned either to clients or to CDN delivery points

Pilot run on Oct. 27th

Process followed

Initial tests in lab environment with simple topology. Simplistic network configuration (e.g., OSPF) for understanding viability of the approach and get experience.

Technology Lab

Pre-production Lab

Integration on a pre-production environment with realistic network topologies, and network configuration as in the production network, with the purpose of assess the solution and solve problems found.

IETF 114

Preparation of ALTO connection for being an element of the production network (server installation, security aspects, flow definition, etc)

Deployment of ALTO as an element of the production network and full integration with TCDN


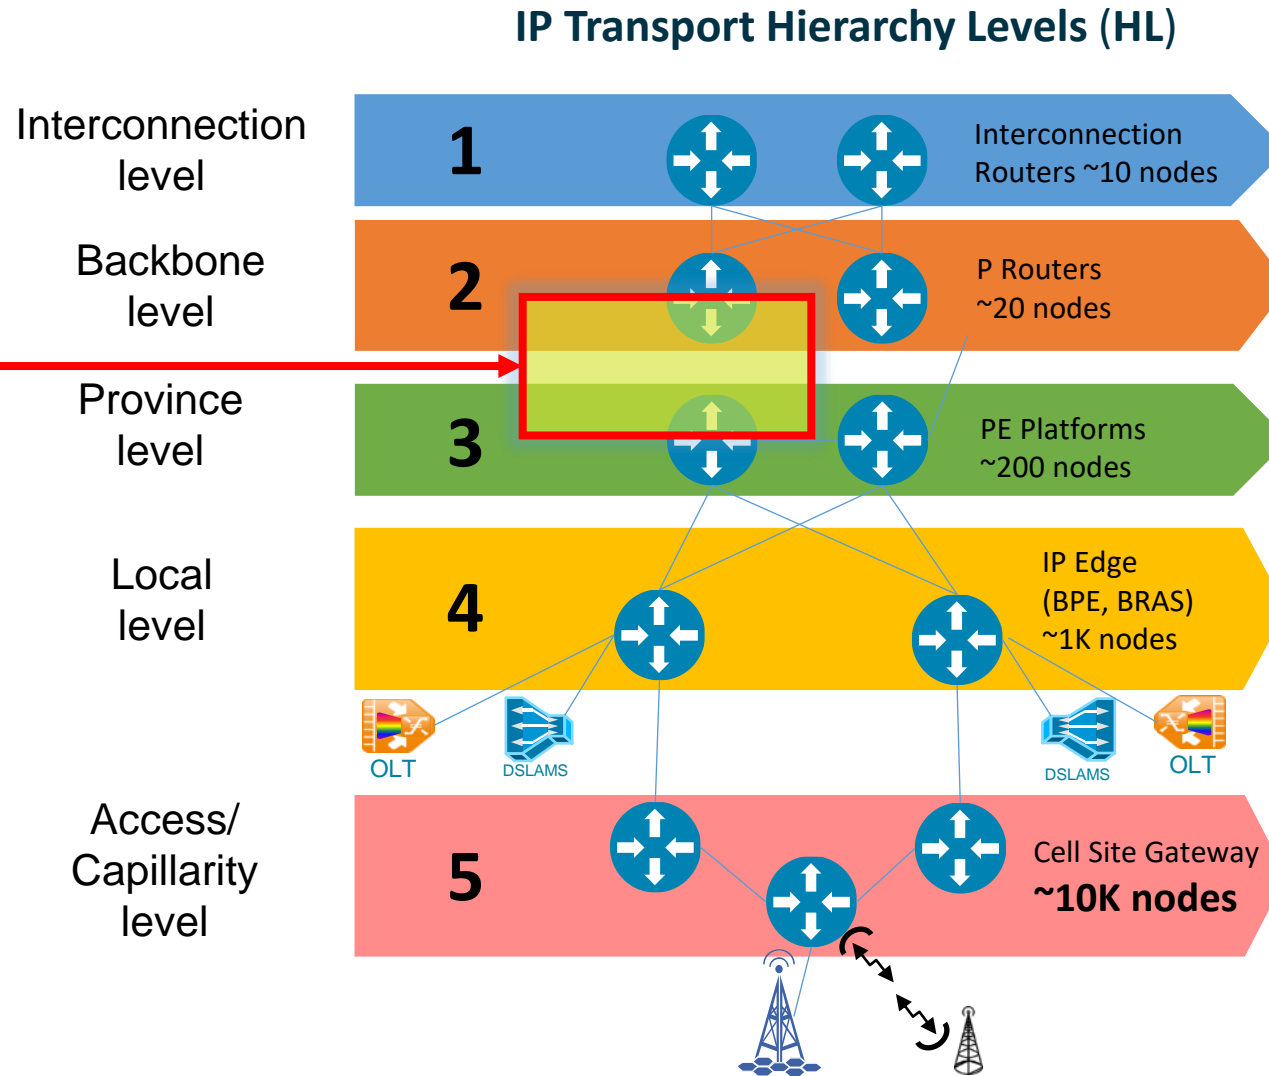
Production Network

IETF 115



Initial known restriction

Due to an issue with current version of the OS of one of the router vendors it is not possible to disseminate BGP-LS information for HL3-HL2 interconnection, except for some few PoPs in a region of the network



High level view of a network in a mid-size country

Bright news

- A total of 16383 summarized IP address ranges are retrieved from the network, allocated to different kind of nodes (fixed, mobile, enterprise)
- Such IP ranges are internal, but also external to Telefonica (those of national interconnections at IXP)
- Available information in cost map correctly reflects the defined IGP metric
- The load of the server is not significant (yet being monitored to understand proper behavior)



Example

Network Map

```
{
...
  "pid0:ac***703": [
    "2.139.150.192/27",
    "2.139.146.32/27",
    "2.138.196.0/27",
    "2.139.170.96/27",
    "2.139.163.0/27",
    "2.139.160.64/27",
    "2.139.153.192/27",
    "2.139.151.160/27"  ],
...
  "pid0:ac***403": [
    "2.139.144.8/29",
    "2.139.146.96/27",
    "2.139.150.160/27",
    "2.139.145.160/27",
    "2.139.148.96/27",
    "2.139.144.32/27",
    "2.139.157.32/27",
    "2.139.159.32/27",
    "2.139.156.224/27",
    "2.139.155.224/27",
    "2.139.151.224/27",
    "2.139.171.160/27",
    "2.139.163.224/27",
    "2.139.165.160/27",
    "2.139.151.96/27",
    "2.139.161.64/27",
    "2.139.164.0/27",
    "2.139.166.0/27" ],
...
}
```

Cost Map

```
"pid0:ac***f01": {
  "pid0:ac***f01": 0,
  "pid0:ac***f02": 1030,
  "pid0:ac***f18": 2530,
  "pid0:ac***f0f": 2530,
  "pid0:ac***f05": 1500,
  "pid0:58***cc7": 5030,
  "pid0:ac***f16": 1500,
  "pid0:ac***f19": 7530,
  "pid0:58***cc3": 10030,
  "pid0:50***3bd": 1300,
  "pid0:50***513": 5000,
  "pid0:c2***293": 1500,
  "pid0:c2***294": 1520
},
"pid0:ac***f02": {
  "pid0:ac***f02": 0,
  "pid0:ac***f01": 1030,
  "pid0:ac***f18": 1500,
  "pid0:ac***f0f": 1500,
  "pid0:ac***f05": 2530,
  "pid0:58***cc7": 4000,
  "pid0:ac***f16": 2530,
  "pid0:ac***f19": 6500,
  "pid0:58***cc3": 9000,
  "pid0:50***3bd": 2330,
  "pid0:50***513": 5000,
  "pid0:c2***293": 1520,
  "pid0:c2***294": 1500
},
}
```

Not so bright news

- There is no information of IP ranges in 4% of the PoPs
 - Analyze what are the particularities of those PoPs
- Some IP ranges seems not to be retrieved
 - check if proper BGP sessions are considered, because of the diversity of RR in the network for different purposes
- Only 27 PIDs are in both network map and cost map
 - Confirm if it is a matter of the HL2-HL3 BGP-LS missing information
- PIDs for CDN nodes not yet captured
 - necessary to connect to another RR?



Next steps

- For the pilot
 - Understand how to consume the ALTO information: how often?
 - Continue analyzing the information received to understand dynamics in a production network
 - Debug issues that could be found during the process
 - Wait till resolution of OS issue for HL2-HL3 connections in the overall network for building a complete picture of the network
- For ALTO-based solution
 - Productification of ALTO
 - Topology load automation to be consumed by TCDN logic
- For ALTO / MOPS WG
 - Document the pilot (would it be also interesting for MOPS?)
 - Identify gaps/issues/improvements in the solution worthy to work in (e.g., security)
 - To provide another update at IETF 116 (with everything solved, hopefully)

Backup


Network map creation

*Leveraging on UPDATE messages
from BGP RR*

Pid name: *Pid + ASN (default value = 0) + IP next-hop in hexadecimal:*

"pid" + 0 + hex(**192.168.255.12**) -> pid0: c0a8ff0c

BGP advertisement with address family ipv4 unicast. The nexthop value (192.168.255.12) is used for the PID identifier and the NLRI value (80.58.102.64/26) is one of the prefixes of such PID



```
{ "exabgp": "5.0.0", "time": 1651759327.1712353, "host" :  
"localhost.localdomain", "pid" : 72259, "ppid" : 1, "counter": 135, "type":  
"update", "neighbor": { "address": { "local": "80.58.171.201", "peer":  
"192.168.252.200" }, "asn": { "local": 64531, "peer": 64531 } , "direction":  
"receive", "message": { "update": { "attribute": { "origin": "igp", "med": 0, "local-  
preference": 100, "originator-id": "192.168.255.12", "cluster-list": [ "3.3.2.1",  
"192.168.251.172" ] }, "announce": { "ipv4 unicast": { "192.168.255.12": [ {  
"nlri": "80.58.102.64/26" } ] } } } } } }
```

```
"pid0:c0a8ff0c": [  
  "17.100.0.1/32",  
  "17.0.0.3/32",  
  "192.168.248.240/30",  
  "80.58.102.64/26",  
  "123.123.123.123/32",  
  "17.0.0.1/32"  
],
```

Cost map creation

Leveraging on UPDATE messages
from BGP-LS RR



bgpls-link messages

```
{ "exabgp": "5.0.0", "time": 1651676896.0184126, "host": "localhost.localdomain", "pid": 70559,
  "ppid": 1, "counter": 733, "type": "update", "neighbor": { "address": { "local": "80.58.171.201",
  "peer": "192.168.255.89" }, "asn": { "local": 64531, "peer": 64531 } , "direction": "receive",
  "message": { "update": { "attribute": { "origin": "igp", "local-preference": 100, "originator-id":
  "192.168.252.178", "cluster-list": [ "12.4.1.1" ], "bgp-ls": { "generic-ls-id-258":
  ["0x00000C0D0000008B"], "igp-metric": 1000 }}, "announce": { "bgp-ls bgp-ls": { "192.168.252.178":
  [ { "ls-nlri-type": "bgpls-link", "l3-routing-topology": 2, "protocol-id": 2, "local-node-descriptors": [ {
  "autonomous-system": 3352 }, { "bgp-ls-identifier": "178" }, { "router-id": "d500b8070000" } ],
  "remote-node-descriptors": [ { "autonomous-system": 3352 }, { "bgp-ls-identifier": "178" }, { "router-
  id": "c0a8ff210000" } ], "interface-addresses": [ "192.168.204.198" ], "neighbor-addresses": [
  "192.168.204.197" ], "multi-topology-ids": [ ], "link-identifiers": [ ] ] ] ] ] } } }
```

bgpls-node messages

```
{ "exabgp": "5.0.0", "time": 1652431307.7986917, "host": "localhost.localdomain", "pid": 81254,
  "ppid": 1, "counter": 774, "type": "update", "neighbor": { "address": { "local": "80.58.171.201",
  "peer": "192.168.255.89" }, "asn": { "local": 64531, "peer": 64531 } , "direction": "receive",
  "message": { "update": { "attribute": { "origin": "igp", "local-preference": 100, "originator-id":
  "192.168.252.178", "cluster-list": [ "12.4.1.1" ], "bgp-ls": { "node-flags": { "O": 1, "T": 0, "E": 0, "B": 1,
  "R": 0, "V": 0, "RSV": 0 }, "area-id": "1346612317385085419520", "local-te-router-ids":
  ["192.168.181.3"] }}, "announce": { "bgp-ls bgp-ls": { "192.168.252.178": [ { "ls-nlri-type": "bgpls-
  node", "l3-routing-topology": 2, "protocol-id": 2, "node-descriptors": [ { "autonomous-system": 3352 },
  { "bgp-ls-identifier": "178" }, { "router-id": "192168181009" } ], "nexthop": "192.168.252.178" } ] ] } } }
```

python
networkx
library



```
{
  "pid0:c0a8fc72": {},
  "pid0:c0a8fc8a": {
    "pid0:c0a8fc8a": 0,
    "pid0:c0a8fca2": 1,
    "pid0:c0a8fcf9": 2,
    "pid0:c0a8fca0": 3,
    "pid0:c0a8fcec": 4,
    "pid0:c0a8fcf5": 4,
    "pid0:c0a8fca3": 4,
    "pid0:c0a8fb7b": 5,
    "pid0:c0a8fcf6": 5,
    "pid0:c0a8fbf1": 5
  },
  "pid0:c0a8fca1": {
    "pid0:c0a8fca1": 0,
    "pid0:c0a8fcf6": 1,
    "pid0:c0a8fcf1": 1,
    "pid0:13d833978": 1,
    "pid0:c0a8fbc8": 1,
    "pid0:c0a8fca3": 2,
    "pid0:c0a8fb7b": 3,
    "pid0:c0a8fca0": 3,
    "pid0:c0a8fbf1": 3,
    "pid0:c0a8fcec": 4,
    "pid0:c0a8fcf9": 4,
    "pid0:c0a8fcf5": 4,
    "pid0:c0a8fca2": 5,
    "pid0:c0a8fc8a": 6
  },
}
```