

CCAMP Working Group
Internet Draft

Y. Lin
Huawei Technologies
L. Han
Y. Zhao
China Mobile
R. Munoz
CTTC
Y. Tan
China Unicom

Category: Informational
Expires: January 8, 2026

July 7, 2025

Applicability of GMPLS and PCEP for fine grain Optical Transport
Network

draft-lin-ccamp-gmpls-fgotn-applicability-03

Abstract

ITU-T Recommendation ITU-T G.709/Y.1331 (2020) Amd.3 (03/2024) introduced new fine grain OTN (fgOTN) for the efficient transmission of sub-1G client signals.

This document reviews the fgOTN control plane requirements, examines the applicability of using existing GMPLS control plane for fgOTN, and provides the standard gap analysis and considerations on GMPLS extensions.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on January 8, 2026.

Copyright Notice

Copyright (c) 2025 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 (<http://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	3
2. G.709 Optical Transport Network Overview	3
2.1. fgOTN Mapping and Multiplexing Architecture	4
2.1.1. fgODUflex(p) into ODUj(fgTS) or ODUflex(fgTS,n) Multiplexing	5
2.1.2. ODUj(fgTS) or ODUflex(fgTS,n) into ODUk Multiplexing	6
2.2. fgOTN Connection Model	6
2.3. fgOTN Use Cases	8
2.3.1. Point-to-Point (P2P) Private Line Service	8
2.3.2. Multi-Point-to-Multi-Point (MP2MP) Cloud Access	9
3. General fgOTN Control Consideration	10
3.1. Signal Type	10
3.2. fgOTN Label	10
4. fgOTN Connection Control Consideration	10
4.1. Connection Hierarchy	10
4.2. Hitless Resizing	10
4.3. Scalability Consideration	11
5. fgOTN Service Control Consideration	11
6. fgOTN Routing Consideration	14
6.1. fgOTN Resource Distribution	14
6.2. Scalability Consideration	14
7. fgOTN Link Management Consideration	15
8. Manageability Considerations	15
9. Security Considerations	15
10. IANA Considerations	15
11. References	15
11.1. Normative References	15
11.2. Informative References	15
12. Contributors	16

13. Authors' Addresses	17
------------------------------	----

1. Introduction

Optical Transport Networks (OTN) is a mainstream layer 1 technology for the transport network. Over the years, it has continued to evolve, to improve its transport functions for the emerging requirements.

OTN control plane capabilities are also introduced and improved according to the evolution of OTN technologies. The OTN control plane brings benefits to operators, including improving OTN network resiliency and resource usage efficiency. Control plane technologies such as Generalized Multi-Protocol Label Switching (GMPLS) [RFC3945] and Path Computation Element Communication Protocol (PCEP) [RFC4655], support different classes of interfaces and switching capabilities including OTN.

In the latest version of OTN, ITU-T G.709/Y.1331 (2020) Amd.3 [G709-202403], the fine grain OTN (fgOTN) is introduced for the efficient transmission of low rate signals (e.g., sub-1G).

In [OCN], the Optical Cloud Network (OCN) architecture is specified, which uses OTN (including fgOTN) as the underlay network to carry the high-quality cloud services. The [OCN] also specifies the functionalities of Optical Service Protocols (OSP), which run on the control interfaces of the OCN for service and connection control.

This document reviews the latest OTN standard technologies (especially the fgOTN) and their requirements to the OTN control plane.

This document also examines the applicability of using existing GMPLS and PCEP for fgOTN, and provides the standard gap analysis and considerations on GMPLS and PCEP extensions.

2. G.709 Optical Transport Network Overview

Recommendation ITU-T G.709/Y.1331 [G709-202403] defines the requirements for the optical transport network (OTN) interface signals of the optical transport network.

The most important new feature introduced by G.709/Y.1331 [G709-202403], compared with the previous editions of G.709 series standards, is the introduction of fine grain OTN (fgOTN) for the

efficient transmission of low rate signals (e.g., sub-1G). This can be an alternative to the existing SDH networks, which phases out in operators' networks.

Recommendation ITU-T G.709.20 [G709.20] provides an overview of functions provided by the fgOTN layer network.

The main functional requirements of fgOTN, described in G.709.20 [G709.20], include:

- Support TDM hard isolation with deterministic latency to guaranteed performance of OTN networks.
- Support mapping packet services and Constant Bit Rate (CBR) services into fgODUflex path layer.
- Support multiplexing multiple fgODUflex signals into ODU0/1/2/flex(fgTS,n) (where n = 3 to 7).
- Support fgODUflex SNCP 1+1 protection.
- Support timing transparent transport for CBR services.
- Support fgODUflex hitless bandwidth adjustment function for packet services.

The detail definition of fgOTN, including frame format, mapping of client signals into fgOTN and mapping of fgOTN into ODUk, is described in G.709 [G709-202403].

2.1. fgOTN Mapping and Multiplexing Architecture

fgOTN layer network is a service layer network of the OTN ODU layer network. Figure 1 shows the overview of fgOTN mapping and multiplexing architecture.

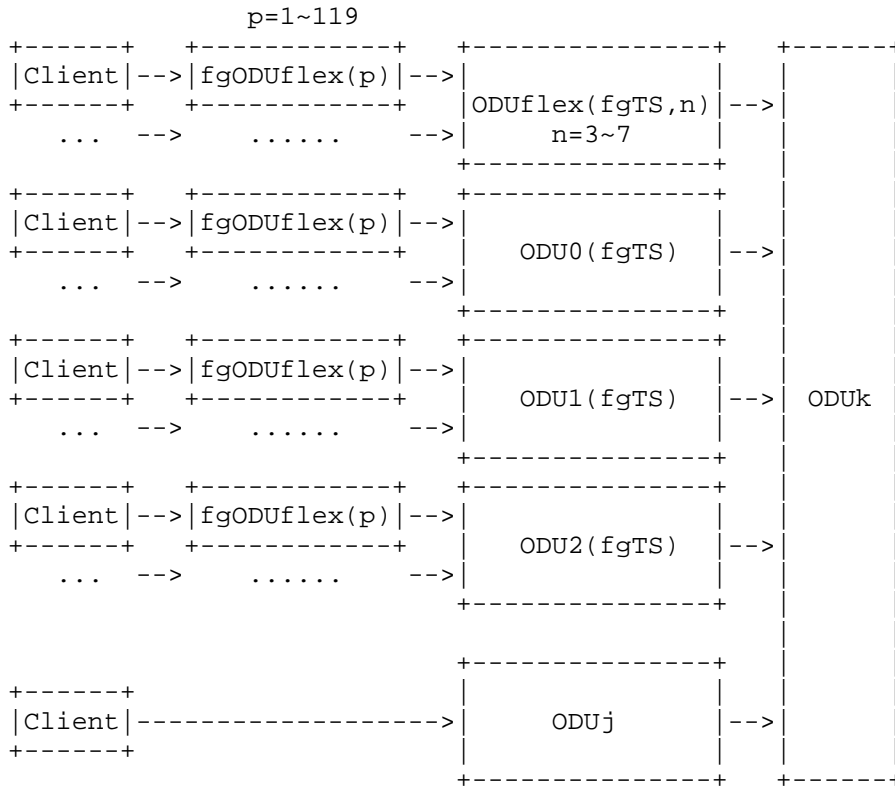


Figure 1: fgOTN mapping and multiplexing architecture

Note that, similar to ODUj into OTUj mapping, the ODUj(fgTS) (j=0,1,2) may also be mapped into OTUj directly, which is not shown in Figure 1 for simplicity.

2.1.1. fgODUflex(p) into ODUj(fgTS) or ODUflex(fgTS,n) Multiplexing

The ODUj(fgTS) (j=0,1,2) and the ODUflex(fgTS,n) are the new containers used for fgODUflex(p), which contain multiple fine grain Tributary Slots (fgTSs) with an approximate bit rate of 10 Mbit/s. More specifically:

- The ODU0(fgTS) contains 119 fine grain Tributary Slots (fgTSs), and the bit rate of each fgTS is approximately 10411 kbit/s.
- The ODU1(fgTS) contains 238 fgTSs, and the bit rate of each fgTS is approximately 10455 kbit/s.

- The ODU2(fgTS) contains 952 fgTSs, and the bit rate of each fgTS is approximately 10499 kbit/s.
- The ODUFlex(fgTS,n) ($3 \leq n \leq 7$) contains $n \cdot 119$ fgTSs, and the bit rate of each fgTS is approximately 10411 kbit/s.

When multiplexing an fgODUFlex(p) into an ODUj(fgTS) ($j=0,1,2$) or an ODUFlex(fgTS,n), the fgODUFlex(p) occupies p fgTSs of the ODUj(fgTS) or ODUFlex(fgTS,n), where p is a positive integer that less than or equal to 119.

Packet or CBR client signals which are less than 1 Gbit/s can be carried by an fgODUFlex(p) to improve the efficiency of OTN bandwidth utilization.

2.1.2. ODUj(fgTS) or ODUFlex(fgTS,n) into ODUk Multiplexing

The ODUj(fgTS) ($j=0,1,2$) or ODUFlex(fgTS,n) can be multiplexed into ODUk using one or multiple 1.25Gbps tributary slots of the ODUk, which is similar to the traditional ODUj (including ODUFlex) into ODUk multiplexing.

Traditional ODUj and ODUj(fgTS) / ODUFlex(fgTS,n) can be co-exist in the same ODUk, i.e., it's possible that some of the 1.25Gbps tributary slots of an ODUk are used for the traditional ODUj, while other 1.25Gbps tributary slots of the ODUk are used for the ODUj(fgTS) / ODUFlex(fgTS,n).

2.2. fgOTN Connection Model

Figure 2 shows a simple example about the fgOTN connection model.

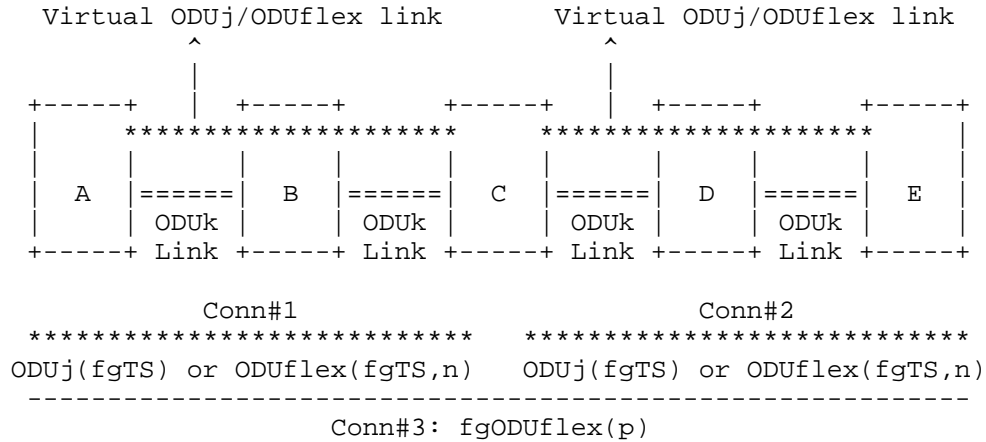


Figure 2: example of fgOTN connection model

In this example, there are five OTN nodes which are interconnected by ODUk links. Two ODUj(fgTS) or ODUflex(fgTS,n) connections are created:

- Conn#1: ODUj(fgTS) or ODUflex(fgTS,n), A-B-C
- Conn#2: ODUj(fgTS) or ODUflex(fgTS,n), C-D-E

These two connections form two virtual links, which are used for the fgODUflex(p) connection.

- Conn#3: fgODUflex(p), A-C-E

From the OTN node point of view (only describe one connection direction for simplicity. The other connection direction is exactly the same):

- Node A multiplexes the fgODUflex(p) into an ODUj(fgTS) or ODUflex(fgTS,n), which are further multiplexed into ODUk.
- Node B and D perform the cross-connect for the ODUj(fgTS) or ODUflex(fgTS,n), without being awareness of the fgODUflex(p) inside.
- Node C demultiplexes the Conn#1 signal from the ODUk, demultiplexes the fgODUflex(p) from the Conn#1, perform fgODUflex(p) cross-connect, and multiplexes the fgODUflex(p) into Conn#2 and then into ODUk.
- Node E demultiplexes the Conn#2 signal from the ODUk, and further demultiplexes the fgODUflex(p) from the Conn#2.

2.3. fgOTN Use Cases

2.3.1. Point-to-Point (P2P) Private Line Service

One of the motivations of fgOTN is to provide an alternative to carry sub-1G private line services. Figure 3 below shows an example of reference network model for multi-domain fgOTN private line service (also see Annex A of G.709.20 [G709.20]).

In this example, there are one backbone OTN domain, two Metro OTN domains and two fgOTN CPEs in the network. The backbone OTN domain only supports ODUk switching, while the Metro OTN networks support both fgODUflex and ODUk switching. Therefore, at the boundary nodes (e.g., metro core nodes) of the metro OTN domains, the fgODUflexes to other metro OTN networks are multiplexed into ODUk of the backbone network. The backbone OTN nodes perform ODUk switching without being aware of the fgODUflexes inside the ODUk.

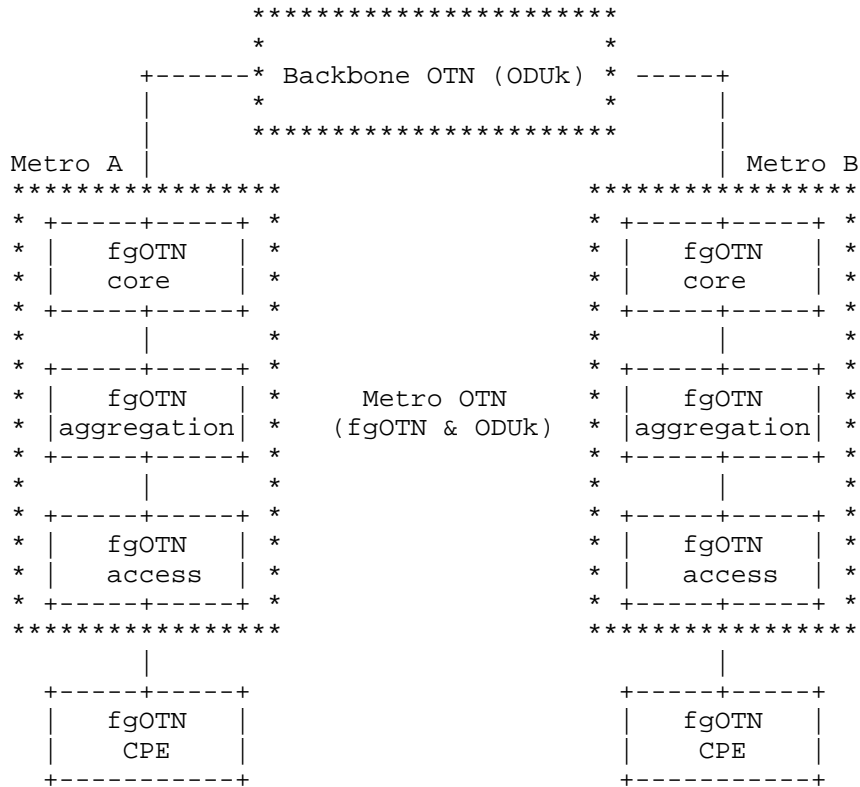


Figure 3: Example of fgOTN-based private line (from [G709.20])

2.3.2. Multi-Point-to-Multi-Point (MP2MP) Cloud Access

Cloud applications are becoming widely deployed in enterprises and vertical industries. Organizations with multiple campuses are interconnected together with the remote cloud Data Centers (DCs) for storage and computing. Such cloud services demand that the underlying network provides Multi-Point-to-Multi-Point connectivity with high quality of experience, such as high availability, low latency, and on-demand bandwidth adjustments. fgOTN, as a TDM-based technology, can be used to meet such requirements. This could be seen as an MP2MP private network service.

[F5G-UC] provides several use cases on how OTN is used for multi cloud access, including enterprise private line connectivity to multiple clouds and premium home broadband connectivity to multiple clouds. To support these use cases, the [OCN] further defines the Optical Cloud Network architecture, and further analyze the technical requirements on the service and connection control protocols.

[opt2cloud] also provides similar multi cloud access use cases, and further provides the problem statement and control plane technical gap analysis on using fgOTN for multi cloud access. See Figure 4 below and see more detail description in [opt2cloud].

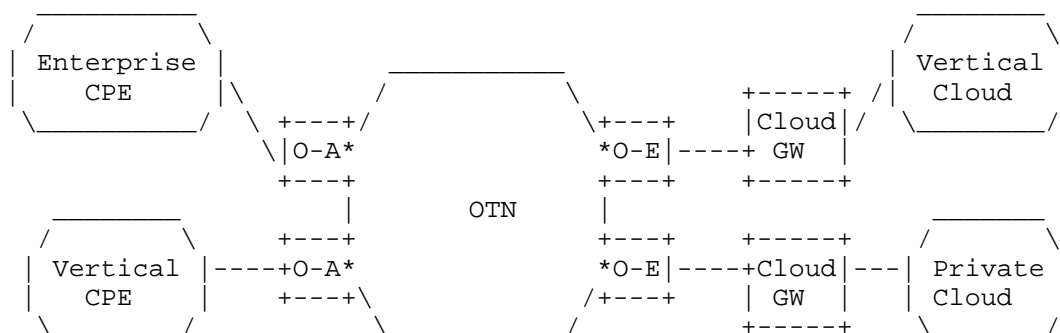


Figure 4: Multi-cloud access through an OTN (from [opt2cloud])

3. General fgOTN Control Consideration

3.1. Signal Type

[G709-202403] introduces new signal types including fgODUflex, ODUj(fgTS) ($j=0,1,2$) and ODUflex(fgTS,n).

The fgOTN control plane needs to support these fgOTN related signal types.

3.2. fgOTN Label

In a control plane for a TDM network, labels are used to indicate the tributary slots in a link which are used by a connection.

In [RFC7139], the ODU label for ODUj into ODUk multiplexing is defined. There are at most 80 tributary slots in an ODUk (i.e., ODU4 link).

In [G709-202403], the bit rate of the tributary slot (10 Mbit/s) is much smaller than before (1.25 Gbit/s), therefore the total number of fgOTN tributary slots will be much larger. There are $119 \cdot n$ ($1 \leq n \leq 8$) tributary slots in an ODUj(fgTS) or ODUflex(fgTS,n) (952 tributary slots in maximum).

An fgODUflex(p) occupies p ($1 \leq p \leq 119$) tributary slots of the ODUj(fgTS) or ODUflex(fgTS,n).

New fgOTN label needs to be designed to support the maximum number of tributary slots.

4. fgOTN Connection Control Consideration

4.1. Connection Hierarchy

As described in Section 2 of this document, two-stage multiplexing may happen when creating an fgODUflex connections (e.g., fgODUflex into ODUflex(fgTS,n) and then into ODUk multiplexing). The fgOTN control plane needs to support the control of the multi-stage multiplexing.

4.2. Hitless Resizing

[RFC7139] supports the control of hitless resizing of ODUflex.

[G709-202403] defines the data plane procedure to support fgODUflex hitless resizing. The support of control of hitless resizing of fgODUflex needs to be further considered.

4.3. Scalability Consideration

fgOTN support much finer granularity of ODU connections. Therefore, the number of ODU connections will be significantly increased.

For example, there could be up to 952 (119*8) fgOTN connections within a single ODU2 link. As a comparison, there are at most 8 ODU0 connections within an ODU2 link and 80 ODU0 connections within an ODU4 link in a traditional OTN network.

Therefore, the scalability and the performance of fgOTN connection control need to be considered. Specifically, when an ODU link fails, there may be thousands of fgOTN connections which are affected by the failure and need to be rerouted at the same time. Even in such case, the performance of fgOTN connection restoration still needs to be guaranteed.

If the General Communication Channel (GCC) overhead bytes are used as the Data Communication Network (DCN) of the OTN, the DCN bandwidth may become the bottleneck when transmitting thousands of rerouting signaling messages at the same time.

A possible way to improve the performance of the connection control is to use the parallel approach, i.e., the connection control messages are sent to each node of the path in parallel. The PCE as a Central Controller (PCECC) ([RFC9050]) may be extended to support parallel fgOTN connection control.

Note that [OCN] provides several fgOTN connection control approaches (including the parallel approach), aiming at improving the performance of connection provisioning and restoration with large number of fgOTN connections.

5. fgOTN Service Control Consideration

Section 2.3.2 of this document describes the scenarios where fgOTN is used for multi-cloud access services. In such scenarios, multiple fgODUflex connections will be created for a customer, to form an MP2MP private network interconnecting multiple customer's branches and multiple cloud DCs.

Since an OTN edge node may have multiple fgOTN connections connected to different destinations for a customer, the OTN edge node needs to support identification of customer's service flows, and support

mapping different service flows (with different destination addresses) into corresponding fgOTN connections. This requires each OTN edge node to be able to:

- Collect the client-side (the customer side or the cloud side, which the OTN edge node are connected to) service address information;
- Note that this step may be done by manual configuration or by running relevant protocols between client side and the OTN edge node, which is out of scope of this document.
- Learn the other client-side service address information from all other OTN edge nodes;
 - Generate the service mapping rules;
 - Identify the service identification information (e.g., source and destination IP or MAC addresses) from the packet headers of the received service flow;
 - Mapping the received service flow into corresponding fgOTN connections, according to the generated service mapping rules.

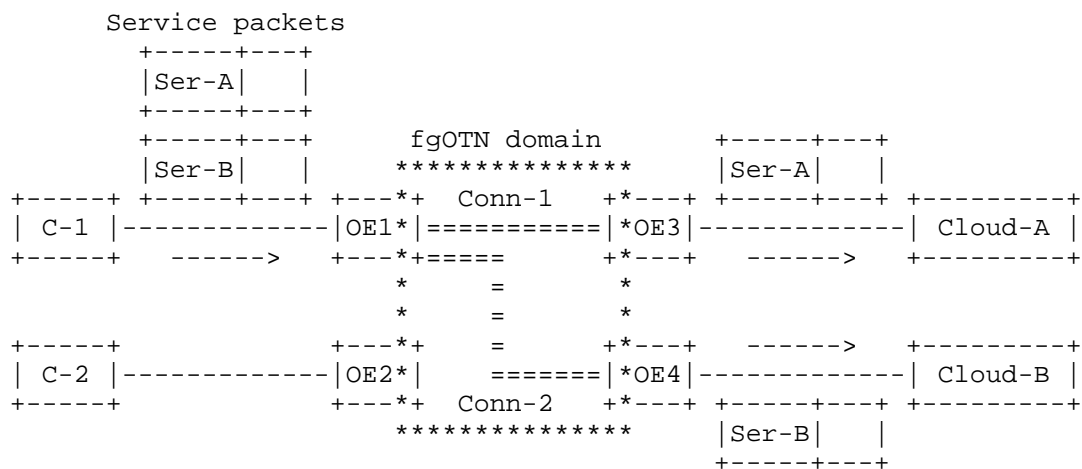


Figure 5: Example of multi-cloud access service

Figure 5 shows an example on multi-cloud access service. In this example, an fgOTN-based private network service is provisioned, with multiple fgOTN connections interconnecting customer's branches (C-1 and C-2) and two cloud DCs (Cloud-A and Cloud-B). When the OTN Edge node OE1 receives the service flow packets with service identifications (e.g., the destination addresses) Ser-A and Ser-B,

OE1 needs to map the service packets identified as "Ser-A" into fgOTN connection #1, and map the service packets identified as "Ser-B" into fgOTN connection #2. In this way, the service flow can be transmitted to the correct destination cloud DCs.

Note that traditional OTN with 1.25 Gbit/s or 2.5 Gbit/s tributary slots can also be used in this scenario.

[opt2cloud] provides more detail description on how the service address information is learned, and how the service flows are mapped to the fgOTN connections.

New protocol in the control plane is needed for the OTN edge nodes, so that they can know which OTN connections can be selected to transmit a certain service flow.

1) Distribution way, with extension to IGP

IGP protocols, such as OSPF-TE and IS-IS-TE, provide a distribution method to flood the topology (including the address information) and resource information within the network.

With future protocol extensions to the OSPF-TE or IS-IS-TE, an OTN edge node is possible to flood the service address information of its client-side network to other OTN edge nodes.

The defect with such flooding method is that, all the intermediate OTN nodes will also receive such service address information. This is not necessary, but brings large amount of extra data transmission and process in the control plane of the intermediate OTN nodes.

2) Centralized way, with extension to PCEP

Another choice would be using a centralized way for the service address learning. Each OTN edge node communicates with a centralized controller via the control plane, to report its client-side service address information. Then the controller reflects such information to other OTN edge nodes directly. This is much efficient compared to the distributed flooding method.

[OCN] introduces the OSP service control procedure, which supports service address learning in a centralized way. In the OSP service control procedure, a centralized Service Mapping Control Component (SMCC) collects the service/client address information from each OTN edge node via a "Service/client node Address Report" message, generates the service mapping information, and sends the service mapping information to each OTN edge node via a "Service Mapping Update" message.

PCEP has been extended to support distribution of link-state and TE information in a centralized way, as defined in [PCE-LS]. With further extensions to [PCE-LS], the PCEP messages could be used as the "Service/client node Address Report" and the "Service Mapping Update" messages in OSP, to support service address learning.

The draft [PCE-fg] is working on the extension to the PCEP to support the service address learning and service flow mapping.

6. fgOTN Routing Consideration

6.1. fgOTN Resource Distribution

[RFC7138] defines the routing protocol for the traditional OTN with 1.25 Gbit/s and 2.5 Gbit/s tributary slot types.

[G709-202403] introduces new type of ODU tributary slot, i.e., the fine grain tributary slot with a bit rate of approximately 10 Mbit/s.

The fgOTN control plane routing protocols need to support the discovery and distribution of fgOTN node and link resource information and capability information, including:

- support of fine grain tributary slot, as well as the fine grain tributary slot resource information;
- fgODUflex multiplexing capabilities;
- fgODUflex switching capabilities.

The typical scenarios for fgOTN is to provide low bit rate private line or private network services for more customers. This implies that the OTN network (which supports fgOTN connections) will cover a larger scope of networks, which may include the backbone network, metro core, metro aggregation, metro access, and even the OTN CPE in the customers' networks.

6.2. Scalability Consideration

Routing protocols such as OSPF-TE and ISIS-TE use flooding mechanisms to distribute the routing information. The route convergence time may increase when the scale of the network becomes larger. Therefore, the scalability of the fgOTN routing protocols needs to be considered.

A large-scale network may be divided into multiple domains, to reduce the route convergence time. Centralized mechanism defined in [PCE-LS] may be used to collect the TE information in a domain by the PCE, to avoid long routing convergence time.

7. fgOTN Link Management Consideration

For further study.

8. Manageability Considerations

For further study.

9. Security Considerations

For further study.

10. IANA Considerations

This document requires no IANA actions.

11. References

11.1. Normative References

None.

11.2. Informative References

[RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.

- [RFC7138] Ceccarelli, D., Ed., Zhang, F., Belotti, S., Rao, R., and J. Drake, "Traffic Engineering Extensions to OSPF for GMPLS Control of Evolving G.709 Optical Transport Networks", RFC 7138, March 2014.
- [RFC7139] Zhang, F., Ed., Zhang, G., Belotti, S., Ceccarelli, D., and K. Pithewan, "GMPLS Signaling Extensions for Control of Evolving G.709 Optical Transport Networks", RFC 7139, March 2014.
- [RFC9050] Li, Z., Peng, S., Negi, M., Zhao, Q. and Zhou, C., "Path Computation Element Communication Protocol (PCEP) Procedures and Extensions for Using the PCE as a Central Controller (PCECC) of LSPs", RFC 9050, July 2021.
- [G709-202403] ITU-T, "Interfaces for the optical transport network", G.709/Y.1331 (2020) Amd.3, 03/2024.
- [G709.20] ITU-T, "Overview of fine grain OTN", G.709.20, 04/2024.
- [F5G-UC] ETSI GR F5G 008 (V1.1.1), "Fifth Generation Fixed Network (F5G); F5G Use Cases Release #2", 2022.06.
- [OCN] ETSI GS F5G 018 (V1.1.1), "Fifth Generation Fixed Network (F5G); Architecture of Optical Cloud Networks", 2024.10.
- [opt2cloud] S. Liu, H. Zhang, A. Guo, Y. Zhao, and D. King, "draft-liu-ccamp-optical2cloud-problem-statement".
- [PCE-LS] D. Dhody, S. Peng, Y. Lee, D. Ceccarelli, A. Wang and G. Mishra, "draft-ietf-pce-pcep-ls".
- [PCE-fg] L. Han, H. Zheng, M. Wang, and Y. Zhao, "draft-han-pce-path-computation-fg-transport".

12. Contributors

Liang Zhong
China Unicom
No. 50, Fuchunjiang East Street, Jianye District
Nanjing City, China
Email: zhongliang1@chinaunicom.cn

13. Authors' Addresses

Yi Lin
Huawei France
Marco Polo A2, 790 Avenue Maurice Donat
Mougins France
Email: yi.lin@huawei.com

Liuyan Han
China Mobile
No.32 Xuanwumen west street
Beijing, 100053
China
Email: hanliuyan@chinamobile.com

Yang Zhao
China Mobile
No.32 Xuanwumen west street
Beijing, 100053
China
Email: zhaoyangyjy@chinamobile.com

Raul Munoz
Centre Tecnologic de Telecomunicacions de Catalunya (CTTC)
Av. Carl Friedrich Gauss, 7 - Building B4
08860 - Castelldefels, Spain
Email: raul.munoz@cttc.es

Yanxia Tan
China Unicom
Email: tanyx11@chinaunicom.cn

