Internet Engineering Task Force

Olivier Duroyon
Rudy Hoebeke
Hans De Neve
Dimitri Papadimitriou
Alcatel

Internet Draft
Document: draft-duroyon-te-ip-optical-01.txt
November 2000
Expiration Date: May 2001

G.LSP Service Model framework in an Optical G-MPLS network
<draft-duroyon-te-ip-optical-01.txt>

Status of this Memo

1. Abstract

The objective of this draft is to propose an IP service model for a non-packet switch capable optical network where G.LSPs are dynamically triggered by the IP layer and subsequently advertised for IP routing. The business model assumes that several IP service domains, some of which represent different administrative entities, share the same optical backbone and focuses therefore primarily on an overlay model. G-MPLS signaling (refer to [g-mpls]) with UNI support is assumed as underlying control plane protocol.

2. Conventions used in this document

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 RFC2119.

3. Introduction

This draft introduces an end-to-end IP service model enabling the dynamic management of Generalized Label Switched Paths (G.LSP) by means of G-MPLS signaling with User-to-Network Interface (UNI) support. A G.LSP is a point-to-point connectivity with specified attributes (some of which are mandatory, while others are optional) established between two termination points in the optical network. A G.LSP could be a fiber switched path, a lambda switched path, a TDM switched path (circuit) or a packet-switch capable G.LSP. The scope of this draft is restricted to optical networks, which are by definition non-packet switch capable. Consequently, G.LSPs are restricted to non-packet switch capable G.LSPs, which we hereafter refer to as G.LSPs.

For reasons of definiteness, the optical devices are always referred to as Optical Network Elements (ONE) and the IP devices as Client Network Elements (CNE). Boundary CNEs and boundary ONEs are interconnected through an UNI signaling and routing interface. The owner of the UNI interface in the optical domain (UNI-Network or UNI-N) is referred to as in the Boundary ONE Controller (ONE-C). Its counterpart in the Client Network (UNI-Client or UNI-C) is referred to the Boundary CNE Controller (CNE-C).

The terminology used in the draft attempts to be in line with the definitions found in [ip-optical], [ouni-framework] and [OIF2000.125.2].

An overlay use of G-MPLS (UNI support) is appropriate for an untrusted environment where several IP service domains, representing different administrative entities, share the same optical backbone. Moreover, this model seems well suited for a network architecture including non-IP devices, e.g., legacy TDM or ATM equipment, that interface with the same optical backbone. This is however beyond the scope of this draft.

To distinguish between trusted and untrusted peers, a separate definition for a Trusted and Untrusted network interfaces is proposed:
- An Untrusted interface is defined when UNI (respectively NNI) interfaces belongs to distinct administrative authorities. For instance an UNI interface between a client network element and an optical network element belonging to distinct ISPs defines an

untrusted relationship between the client and the optical network element.
- A Trusted interface is defined when UNI (respectively NNI) interfaces belongs to the same administrative authority. For instance an NNI interface between two ONEs belonging to the same

optical carrier defines a trusted relationship between these optical networks elements.

The service model further assumes that a decision point in the IP service domain, e.g., a Traffic Engineering tool (TE tool), will trigger a boundary CNE to issue a G.LSP request towards the optical domain. The decision point determines the need for a G.LSP on the basis of IP Service Level Agreements (IP SLAs) and related information, such as for instance load measurements in the IP service domain.

The same TE-tool may also decide about the configuration of Traffic Engineering LSPs (TE-LSP), which are by definition Packet-switch capable LSPs. For the purpose of IP traffic engineering, TE-LSPs are in this case created on top of non-Packet-switch capable G.LSPs.

In the remainder of the document, the terms TE tool or decision point are used interchangeably and refer to the IP service domain device, capable of triggering G.LSP requests.

4. IP service model description

This section outlines the sequence of events that characterize our proposed IP service model.

(1) Configuration

Configuration consists of installing and configuring interfaces of the boundary ONEs and boundary CNEs.

During this stage of end-points configuration, physical attributes of the end-point (as protection attributes of the drop side) are also configured. Configuration of the NNI interfaces of the ONEs is out of the scope of this draft.

(2) Neighbor Discovery, Endpoint Registration and Service Discovery

The objective of Neighbor Discovery is to provide the information needed to identify the neighbor identity and neighbor connectivity over each link interconnecting a boundary CNE to a boundary ONE.

Endpoint registration is concerned with registering boundary CNE endpoints to the optical network. The registration information includes the resource capabilities, closed user group (CUG) identification, port reachability information, UNI protection capabilities etc. This set of information is critical to enable dynamic G.LSP services at the UNI. The endpoint registration mechanism enables the end system to register its critical set of information so that other end systems can identify its existence and network properties.

Service discovery is concerned with obtaining the essential information about services from the attached optical network that are available for the CNE. This information is used by CNEs to establish the service environment. The service discovery mechanism allows the network element to convey information about available services to the end system.

After finishing neighbor discovery, endpoint registration and service discovery, each end system should establish the service environment and have sufficient information to generate G.LSP service request. These mechanisms complement each other and they do not depend on the establishment and use of signaling channels between the two parties.

(3) Optical Service Level Agreements

Next, the service model consists of negotiating Optical SLAs (O-SLA) at optical network-client network boundaries, or between optical networks.

In case of an untrusted peering relationship (i.e. untrusted UNI, respectively NNI), each G.LSP request is authenticated and validated against the O-SLA. The validation process against an O-SLA includes checking whether the request is conforming to the restrictions (e.g., on scope) defined in the O-SLA.

O-SLAs may also be defined at trusted interfaces as the optical domain to provision resources that could use them. Trusted in this context refers to the fact that you don't expect the CNE to violate this O-SLA, and as such requests received from trusted neighbors don't need to be validated against the O-SLA.

(4) G.LSP service request

The decision point of the client network determines the required connectivity through the optical domain based on service

requirements as per the IP SLAs. It then triggers the boundary CNEs to send a G.LSP service request towards the associated boundary ONEs, using G-MPLS signaling with UNI support. This process is dynamic and may involve, amongst others, the creation of additional G.LSPs, the deletion of existing G.LSPs or the modification of existing G.LSPs.

(5) Address resolution

At the UNI, the G.LSP service request sent by the CNE needs to include the ONE source and destination termination-point identifiers (in case of trusted UNI interface) or the CNE source and destination termination-point identifiers (in case of untrusted UNI interface). CNE termination-points should also be considered when the G.LSP is established through several optical networks belonging to different administrative authorities.

Consequently, the source client needs to send an address resolution request to obtain the ONE destination termination-point ID or CNE termination-point ID corresponding to the CNE destination logical-address of the G.LSP service request.

(6) Optical path selection

In case a dedicated instance of an IGP is used inside the optical transport network, each boundary ONE learns the complete topology of the optical domain. A Constraint-based Shortest Path First (CSPF) algorithm can then be implemented in the boundary ONEs to calculate a route for the G.LSP in line with the constraints specified in the request. As an example, the route of a G.LSP may depend on the protection requirements or routing constraints specified in the G.LSP request. The latter may indicate that the requested G.LSP should be routed diversely from other G.LSPs. This CSPF algorithm is expected to be quite different from an IP CSPF algorithm because of optical networking specific considerations.

(7) G.LSP advertisement to the IP layer

As soon as the G.LSPs are lit up, they are advertised to the client network. The involved boundary CNEs will either create a new IP link and start to exchange routing information (using IGP or eBGP) or modify the characteristics of the existing IP link.

(8) Traffic engineering for optimization of the optical domain

Optionally, the optical domain may have its own off-line Optical Traffic Engineering (O-TE) tool. This tool may be used for

optimization of resource utilization in the optical network by rearranging some G.LSPs.

5. UNI discovery and registration services

In order to provide a flexible and end-to-end IP Service model, with a minimum set of local provisioning, specific mechanisms and procedures have to be defined at the boundary between the client and the optical network:
- to discover neighbors identity and connectivity
- to register client end-point identity
- and to discover the supported UNI and network services.

Transport mechanisms used for the UNI discovery and registration services are referenced in [OIF2000.125.2] and [OIF2000.200].

5.1 Neighbor discovery at the UNI

The key objective of Neighbor Discovery at the UNI is to provide the information needed to identify the neighbor identity (IP address associated to the corresponding UNI) and neighbor

connectivity over each link connecting the boundary CNE to the boundary ONE.

Neighbor discovery process which is also referred to as the Termination-port identity process, provides the following information to the boundary CNE and ONE:
- the ONE discovers the identity of the client CNE by automatically discovering the IPv4 address assigned to the UNI-C and the identity of each physical port connected to the CNE
- the CNE discovers the identity of to the connected ONE by automatically discovering the IPv4 address assigned to the UNI-N and the identity of each physical port connected to the ONE

If the signaling transport mechanism is not explicitly configured, the neighbor discovery process ends by the bootstrapping of the signaling control-channel used to exchange the information during the end-point registration and the service discovery processes.

5.2 End-point Registration and UNI Service Discovery

The end-point registration process includes the exchange of information between the CNE and ONE for each of the ports and logical ports connecting the CNE to the ONE. A logical port defines the structure of a physical port by identifying for a given port each of the channels included within this port and sub-channels included within this channel.

The UNI Service discovery process includes the exchange of resource-related information of the Framing and Bandwidth capacity of each of the ports and logical ports connecting the CNE to the ONE. Additional parameters, such as the UNI drop-side protection attributes (UNI client-side protection and UNI network-side protection) and the G.LSP Directionality support could also be exchanged during the resource discovery process. For SDH/Sonet interfaces, the Transparency levels (STE-C, LTE-C), the client support of Virtual Concatenation (VC) and the levels of Continuous Concatenation (CC).

The end-point registration process includes also the address registration process [OIF2000.261.1]. The address registration process allows the CNE to register the CNE logical addresses attached to the CNE Termination-point ID to which corresponds an unique ONE Termination-point IDs. A CNE Termination-point ID includes the unique IP address associated with the client element and the logical-port ID. The logical-port ID comprises the port-ID, Channel-ID and Sub-channel-ID as defined in [OIF2000.125.2].

When the address registration is part of the end-point registration process, the CNE associates the CNE Termination-point ID with the corresponding logical address and ONE termination-point ID. When the CNE does not associate logical addresses with their interfaces,

the CNE termination-point ID resolution implies that the boundary ONE knows the mappings between the CNE termination-point ID and the ONE termination-point ID. This case is considered as a particular case where the CNE logical address fields are empty. In this case, the value of the logical address could correspond to the user-group identifier to which the G.LSP belongs; however, in this particular case, the address resolution is always based on the CNE termination-point ID.

Other client identifiers could be exchanged during end-point registration process:

- the CNE registers the Contract ID attached to a specific element and/or interface
- the CNE registers the Closed User-Group (CUG) IDs (i.e. User-Group ID) attached to a specific client end-point or port

5.3 Network Service Discovery

The network service Discovery consists of the G.LSP service-related discovery process and a policy related service discovery process.

During the G.LSP-related service discovery process, the CNE registers and/or discovers the parameters related to

- SDH/Sonet related services, i.e., the SDH/Sonet Transparency levels supported and the Continuous Concatenation levels supported
- G.LSP Directionality support
- Network-side Protection, i.e., the Protection-levels services provided by the internal optical network (Unprotected, Dedicated 1+1 Protection, Dedicated 1:1 and Shared Protection)
- G.LSP Priority classes and Preemption levels supported by the optical network
- G.LSP Diversity options supported by the optical network
- Security levels support (IPSec or other authentication mechanism) within the signaling used on the control-plane throughout the optical network

The discovery of the Policy-related service may include the following parameters:

- Service-levels offered by optical network
- Scheduling-related service supported by the optical transport network and/or the scheduling desired by the client
- Billing-related service supported by the optical transport network and/or the billing method desired by the client
- Vendor-related and Optional parameters

6. Address resolution

As stated in section 4.5, the source client needs to send an address resolution request to obtain the ONE destination termination-point ID (trusted UNI interface) or CNE termination-

point ID (untrusted UNI interface) corresponding to the destination CNE logical-address.

Consequently, at a trusted UNI interface, the G.LSP create message sent by the CNE to the ONE includes the source and destination ONE (or CNE) termination-point IDs requested for this G.LSP. This implies that the source ONE must perform a internal address-lookup toward a directory service or a local mapping table, in order to find the mapping between the destination CNE termination-point ID and the destination ONE termination-point ID.

So, the optical network client only needs to know the CNE source and destination logical address and termination-point ID in order to request a G.LSP creation; any other topological information concerning the optical network termination-point identification is

transparent for the client.

This mechanism is also adapted for inter-domain G.LSP (cf. Section 9.1) since in this case only the autonomous-system (AS) boundary ONE termination-point to CNE termination-point mapping-list has to announced to the neighboring BGP AS's.

7. Optical Service Level Agreements

An optical domain-IP service domain boundary coincides with a UNI with its associated O-SLA. Similarly, if there are multiple optical sub-networks in the optical domain, there will be O-SLAs negotiated at each optical sub-network boundary. An optical sub-network boundary then corresponds to an optical Network-to-Network Interface (NNI). In this draft, we limit the discussion to O-SLAs at the level of UNIs.

As mentioned before, G.LSP requests issued by a boundary CNE are only accepted within the constraints of an O-SLA. This means that in case of an untrusted peering relationship, each G.LSP request is authenticated and validated against the O-SLA. It was already indicated that O-SLAs may also be defined at trusted interfaces. However, G.LSP requests received from trusted neighbors don't need to be validated against the O-SLA.

In the scope of this draft, we only discuss the technical aspects of an O-SLA. Borrowing from the terminology introduced in [diffserv-framework], we refer to the technical part of an O-SLA as an Optical Service Level Specification (O-SLS). An O-SLS is considered to be unidirectional and to specify performance expectations (i.e., the service level) for the IP service domain as well as imposed reachability constraints, e.g., CUG.

O-SLS parameters could for example include:

1. Capacity constraints

An ingress O-SLS may contain limits on the maximum number of G.LSPs that can be established from a specific ingress point, possibly as a function of time of day, as well as bandwidth constraints (OC-48, OC-192, etc.).
An egress O-SLS may put capacity constraints on the G.LSPs that the receiving IP service domain is willing to terminate.

2. Service performance parameters

Examples are G.LSP setup and/or recovery admitted latency,

supported protection/restoration options, availability, supported routing constraints, accessibility (i.e., G.LSP request blocking probability), responsiveness (specifying upper limits on the processing time of G.LSP requests), etc.

3. Constraints on the 'scope' of G.LSP request

This may be viewed as an extension to the concept of CUGs, which by nature already exhibit reachability limitations. Scope constraints are intended to additionally restrict the topological extent of G.LSPs. In its simplest form, the O-SLS offers to accept any G.LSP request issued by the IP service domain over a specific O-UNI up to a maximum capacity without any scope constraint within the CUG (so-called hose O-SLS). Conversely, the agreement may be constrained by the egress point of a G.LSP. For example, the optical domain service provider might agree to the setup of G.LSPs, up to a certain maximum capacity, but only if these G.LSPs are destined to a specific set of egress points within the CUG.

Part of the purpose of O-SLSs is to protect resources in the optical domain by validation of submitted G.LSP requests. If the optical domain and the IP service domain are under control of the same administrative authority, then there is likely to be a trusted peering relationship between both domains. Conversely, in case of an untrusted peering relationship, the optical domain service provider validates incoming G.LSP requests as per the O-SLS. This validation process can be implemented in the ONE-C. In this case, there exist several mechanisms to install an O-SLS in an ONE-C, e.g., CLI, SNMP, LDAP or COPS. Alternatively, the O-SLS enforcement may be outsourced to another policy entity.

An O-SLS offers to accept G.LSP requests at the service level agreed with the IP service domain. The optical domain service provider will provision the optical domain accordingly. A broad range of optical services could be envisioned. As an example, services could be defined with different levels of accessibility, depending on the probability that a G.LSP establishment request is blocked. Moreover, services could also be categorized as protected or non-protected, depending on the offered protection level. All of these service level characteristics influence the resource provisioning process in the optical backbone.

For each G.LSP request, the optical domain service provider may also need to identify the O-SLS for which the request is submitted. Some authentication may be included in the request in order to verify that the rightful IP service provider issued the request. In

some cases, this customer might be implicitly derived from the signaling channel on which the G.LSP request was received. The authentication mechanism must be specified in the O-SLS.

Although it can be assumed that O-SLSs will be static for the foreseeable future, this draft does not preclude dynamic O-SLSs. These would necessitate some automated form of interaction between the IP service domain and the optical domain. In case of an O-SLS at the O-UNI, this may for instance require the interaction between a Bandwidth Broker (BB) in the IP service domain and a Lambda Broker (LB) in the optical domain. At the level of an O-NNI, this would be between different LBs, acting on behalf of the different optical sub-networks. This automated (re-)negotiation of O-SLSs would in turn call for an automated O-SLS admission control function. Note that this admission control function is different from the validation of G.LSP requests as per the negotiated O-SLS, referred to as O-SLS enforcement.

8. G.LSP triggers

As stated in [ip-optical], the G.LSP request triggering process should be part of a stable traffic engineering tool in the IP service domain as opposed to a data-driven shortcut approach, similar to the schemes proposed for IP over ATM networks. Henceforth, an integrated TE-LSP and G.LSP triggering process is proposed at the end of this section to alleviate the shortcomings of the former method and is elaborated below.

8.1 Data-driven shortcut approach for G.LSPs

The data-driven shortcut model would imply that the boundary CNEs use traffic measurements to autonomously control the number of G.LSPs that connect them with a set of remote boundary CNEs across the optical domain. For example, boundary CNE A could detect that some of its traffic is reaching boundary CNE B in a multi-hop way. If this traffic trunk is large enough, boundary CNE A might decide to set-up a G.LSP to boundary CNE B, relieving the IP forwarding at all intermediate CNEs on the multi-hop path. In an overlay model, once a G.LSP has been established to a new destination, it should be announced as a (new) IP link in the IP service domain routing database. As such, it can be used by any TE entity in the IP service domain and this IP link may carry several TE-LSPs. This implies that the TE entity in the IP service domain would then be constantly reacting to decisions of the boundary CNEs that are continuously changing the IP topology.

Such a layered scheme of G.LSP requests and TE-LSP requests is inefficient and could also break the TE service model, when the

only available G.LSP between two boundary CNEs would be torn down. Such a decision might be based on the valid observation that the traffic pattern has changed such that the existing G.LSP is under-utilized and may be re-directed towards another boundary CNE. However, the G.LSP might still carry TE-LSPs. Turning off the G.LSP has the effect of a link failure and will hence trigger the TE entity in the IP service domain to recover from this failure. Depending on whether the TE-LSP was protected or not, one of the following scenarios will take place.

8.1.1 Protected TE-LSPs

TE-LSPs can be used to carry mission critical traffic requiring a fast recovery scheme in case of link failures. Upon such an event, the traffic of the working TE-LSP can be switched to a protect TE-LSP that has been pre-configured along a node- and link-disjoint path. Depending on whether G.LSP is protected or not throughout the optical network, the following alternative is considered:

- Protected G.LSP: if the turned-off G.LSP was protected within the optical domain, the TE-LSP path calculation might have selected this IP link for both the working and the protect path of the TE-LSP. In that case, the TE-LSP protect path will not be available and connectivity will be lost.

- Unprotected G.LSP: in this case the problem would not arise since the route diversity TE-LSP protect scheme would have selected another IP link for the protect path.

8.1.2 Unprotected TE-LSPs

If the TE-LSP was not protected, the source nodes of the TE-LSPs running over the turned-off G.LSP will start a CSPF calculation to find an alternative path. As all source nodes will be competing for the same resources, some G.LSP requests will be blocked and it might take a while before all G.LSPs have been restored.

The above scenario equally pertains to the case of any link failure in an IP service domain. However, link failures in an IP service domain may be considered as rare events. This is however not the case when this link failure behavior is the result of a data-driven shortcut approach across an optical backbone.

8.2 Integrated TE-LSP and G.LSP triggering process

Given the above shortcoming, boundary CNEs should not autonomously decide to tear down a G.LSP. Yet, it may not always be appropriate to maintain an under-utilized G.LSP. However, a G.LSP should not be turned off until the TE-LSPs it carries, have been re-routed along an alternative path. This might even require an additional G.LSP

setup between two other boundary CNEs. All of this calls for a coordinated TE-LSP and G.LSP triggering process, integrated in the

same decision point. This is possible since both responsibilities reside within the IP service domain.

The ability to dynamically establish G.LSPs adds an extra dimension to the TE capabilities of an IP service domain. In addition to forwarding packets along non-shortest paths, it is now also possible to (re-)configure the topology of the IP service domain by means of adding, deleting or modifying G.LSPs across the optical backbone.

This integrated decision point will use the set of IP SLAs and the derived traffic trunk requirements across the IP service domain (possibly complemented with traffic measurements) to determine the optimal set of G.LSPs and TE-LSPs.

Several setup optimization strategies are possible depending on the business model in use between the IP Service domain and the optical domain, and also the assumptions taken on the pre-existing optical topology.

The TE decision point has the complete knowledge of the IP Topology, all optical end-points, including their logical, and physical attributes (granularity, protection attributes, _).

The different strategies may be chosen among the following:

1- Minimize the number of G.LSPs to be lit up

This strategy fits in business models where the optical domain doesn't belong to the service domain, and as such each additional network G.LSP is an additional cost to the service domain. The TE decision point optimizes the number of G.LSPs to set up through the optical domain for a given IP traffic pattern.

2- Add capacity without rearranging optical topology

Before triggering new G.LSPs, the TE decision point tries to rearrange TE-LSPs without rearranging the underlying optical topology.

3- Add capacity with specific explicit constraints

Some environment may lead to some specific constraints to be taken into account during route computation.

One simple example is a mixed ATM / IP network. In this example TE-LSP used by ATM and their underlying G.LSP must not be rearranged during the computation to add optical capacity. The TE decision point optimizes the number of G.LSP (and subsequently TE-LSP topology) with the possibility of pinning down some components (TE-LSP, G.LSP, _)

4- Optimize IP topology without any optical constraint

TE decision point optimizes the IP topology without taking any constraint on number of G.LSPs setup. The only constraints taken are in this case coming from the end-points attributes.

In addition to the computation algorithm strategy, the TE decision point also takes into account the sort of IP services to be achieved, in order to achieve a consistent restoration between protocol layers.

One simple way is to define a linear hierarchy between IP services.

1.- Layer 1 protection - Non-PSC Level Protection

This service only applies for IP link built between two PSC-capable end-points. The G.LSP connecting both end-points is totally protected. It means that it will be chosen from a pool of G.LSPs with source and destination drop-side protection (1+1, 1:1, Shared Protection). And in addition the G.LSP will request a network-protected path via the optical network.

This service will be mainly seen in a traditional environment where the service domain lies on a very reliable transport layer, dedicated to any fast restoration mechanism.

2.- Diverse Layer 2 _ PSC Level Protection

This service also only applies for a G.LSP built between two PSC-capable end-points (for instance, an IP link connection). Two G.LSPs are requested to the optical cloud via the same CNE-to-ONE interface, using source and destination drop-side protected G.LSPs.
No optical or SDH/Sonet network protection are required for the G.LSPs. But diverse optical paths are requested for both G.LSPs.

This service makes sense in a network architecture where the CNE is locally connected to an ONE, and so the diverse path

routing must start at the first ONE of the optical network.

3. Diverse Layer 2 - Network Level Protection

This service applies indifferently in a mixed PSC-capable (and particularly for IP services) and non-PSC capable optical environment, and not necessarily at the boundary CNE.
Two G.LSPs are requested from the IP service domain using two diverse paths. In this case when the G.LSP request reaches the optical cloud boundary, there is no specific protection requirements towards the optical cloud.

4. No G.LSP protection

This service applies when the restoration mechanisms don't rely on pre-existing backup paths. In this case on protection constraints have to be taken into account at the optical layer.

As described in this paragraph, in order to create a consistent end-to-end IP Service Model, network devices and TE decision point have to synchronize each other to setup and maintain the adequate and optimal set of G.LSPs and TE-LSPs. The resulting topology is based on IP services requirements (Protection scheme, _) and computation strategies (Business models, _).

This leads to needs for potential new standardization items, as information exchange between routers and TE decision point (in case of G.LSP setup failure, _). This will be tackle via subsequent studies.

9. G.LSP advertisement to the IP layer

The decision point may thus trigger the set-up of an additional G.LSP to an already connected boundary CNE. Alternatively, it may trigger a rearrangement of existing G.LSPs, or the establishment of a G.LSP to a boundary CNE that could previously not be reached through the optical domain. It might very well be that the decision point triggers a boundary CNE to drop a G.LSP if its capacity is no longer needed to meet the requirements of the IP SLAs.

In order to make efficient use of the dynamicity of the G.LSP create requests, the routing protocol parameters should be dynamically configurable as well. This section outlines a proposal to achieve a seamless integration of a new G.LSP within the IP service domain for the overlay model by means of automatic configuration.

As soon as a G.LSP to a particular boundary CNE has been lit up, it is assumed that it is promoted to an operational IP link. In case of, e.g., regular SDH/SONET framing, this is achieved by running PPP protocol on the newly established G.LSP.

9.1 G.LSP set-up to a previously unreachable boundary CNE

The draft [ompls-ospf] defines the different facets of the creation of an IP link in case of a peer-to-peer model and proposes to use the newly established IP link as a forwarding adjacency in the IP service domain.

The overlay model imposes different requirements on the IP layer of the boundary CNEs. Indeed, once the first G.LSP is established between two boundary CNEs and promoted as IP link, it is to be advertised as a point-to-point link for IP routing in order to

initiate the IP connectivity between the two boundary CNEs. And subsequently it will allow IP reachability between the associated IP service domains.

Two cases must be considered. The G.LSP is promoted to an IP link connecting:

- two boundary CNEs of the same Autonomous System (IGP peering), or,

- two boundary CNEs of different Autonomous Systems (eBGP peering).

The IGP and BGP peering cases do not require the same kind of configuration and are described separately.

Note that in case of an IGP peer, it is necessary that the G.LSP be bi-directional, because IGP protocols require a bi-directional transport layer. Bi-directional G.LSP setup is further detailed within [g-mpls] and [onni-framework].

From an addressing point of view, the Packet switch capable end-points can be unnumbered (and, e.g., identified by the Router Id of the boundary CNE), or numbered through initial configuration, but different from the IP address assigned to the UNI signaling agents (UNI-Client and UNI-Network) terminating the signaling channel.

It has to be noticed again that within the overlay model, the signaling channel identification is neither known nor advertised throughout the IP service domain.

9.1.1 IGP support

Once the first IP link is established between two boundary CNEs and configured to support an IGP peer, the boundary CNEs need to get the proper information:

- The first requirement is to select IS-IS or OSPF for the newly formed IP link.

- In addition, link routing parameters such as timers and area numbers might have to be specified. For instance, timers in OSPF should be consistent at both ends of the IP link.

- Also, link metrics (e.g., resource classes, etc.) need to be inherited or configured for use by IP routing.

- Finally, the IGP protocol is enabled and the IP link is advertised.

These parameters have to be accessible and are automatically configured (prior to or at the time of G.LSP establishment) by the

decision point of the IP service domain to efficiently deploy the IP service on top of the G.LSP.

However, some of the routing parameters (e.g., OSPF timers) may be defaulted to pre-determined values. Those values must be defined network-wide and be consistent between all possible boundary CNE pairs. The decision point should be allowed to overwrite those parameters at the setup time of the G.LSP.

9.1.2 eBGP peering

In the case of an inter-domain G.LSP, static route configuration specifying the BGP peer (most probably a virtual interface of the remote boundary CNE) should be configured in the local boundary CNE in order to set up the TCP session used in BGP.

In addition, an IP SLA is going to be negotiated between the autonomous systems and routing policies are going to be configured on both ends of the G.LSPs.

As opposed to the IGP peering case, triggering of inter-domain G.LSPs will very likely not arise from an automated process because of the BGP peering negotiation procedure.

9.2 Set-up of an additional G.LSP

In addition to the option described in 9.1 (creation of a new stand-alone IP link with the new G.LSP and advertisement to the routing protocol), a second option is now possible, which is to create an additional G.LSP to an existing IP link that forms a bundled link [mpls-bundle]. In this case there is no new configuration necessary for the IGP or BGP routing layer but only interactions internal to the boundary CNEs.

This bundle is advertised as a single IP link. The G.LSP in itself may be unidirectional, and hence the bundle could have an asymmetric bandwidth.

- The boundary CNE upgrades the bandwidth of the bundle link based on the characteristics of this new G.LSP.

- The new G.LSP is included in the load balancing mechanism that distributes the traffic amongst the component G.LSPs of the bundled link, e.g., proportional to their bandwidth.

- The addition of a new G.LSP to a bundle does not impact the routing topology. Only the new bandwidth of the IP link is advertised within the IP service domain. The other characteristics of the IP link, e.g., the resource classes, remain unchanged.

In the case described in this section, the only mandatory information to be automatically configured by the decision point is the bundle identifier to which the G.LSP is to be added.

9.3 Rearrangement of an existing G.LSP

Within the optical network, two alternatives could be considered for the rearrangement of an existing G.LSP:
- Either the non-destructive modification of an already established G.LSP. In this case, the source and destination termination-points of the G.LSP can not be changed, but other parameters such as bandwidth and network protection could be modified without disrupting the working G.LSP.
- Or the destructive modification of an already established G.LSP. This case is the straightforward combination of G.LSP tear-down followed by a new G.LSP set-up towards a different destination.

10. References

[diffserv-framework] Y.Bernet et al, "A Framework for Differentiated Services", Internet Draft, draft-ietf-diffserv-framework-03.txt, February 1999.

[g-mpls] Peter Ashwood-Smith et al., "Generalized MPLS _ Signaling Functional Description", Internet draft, draft-ietf-mpls-generalized-signaling-00.txt, October 2000.

[ip-optical] James Luciani et al., " IP over Optical Networks _ A Framework", Internet draft, draft-ip-optical-framework-00.txt, February 2000.

[ompls-isis] K. Kompella et al., "IS-IS Extensions in Support of MPL(ambda)S", Internet Draft, draft-kompella-isis-ompls-extensions-00.txt, July 2000.

[ompls-ospf] K. Kompella et al., "OSPF Extensions in Support of MPL(ambda)S", Internet Draft, draft-kompella-ospf-ompls-extensions-00.txt, July 2000.

[mpls-bundle] K. Kompella et al., "Link Bundling in MPLS Traffic Engineering", Internet Draft, draft-kompella-mpls-bundle-02.txt, July 2000.

[onni-framework] D.Papadimitriou et al,. "Optical NNI Framework and Signaling Requirements", Work in progress, draft-papadimitriou-onni-framework-00.txt, November 2000.

[ouni-framework] B.Rajagopalan et al., `Signaling Requirements at the Optical UNI', Internet Draft, draft-bala-mpls-optical-uni-signaling-00.txt, July 2000.

[OIF2000.125.2] B. Rajagopalan et al., "User Network Interface v1.0 Proposal", OIF Contribution 125.2, October 2000.

[OIF2000.200] D. Pendarakis et al., "Signaling Transport Mechanisms for UNI 1.0", OIF Contribution 200, September 2000.

[OIF2000.261.1] D. Papadimitriou et al., "Address Registration and Resolution", OIF Contribution 261, November 2000.

11. Acknowledgments

The authors would like to thank Emmanuel Desmet, Sitaram Kalipatnapu and Gert Grammel for their constructive comments.

12. Author's Addresses

Olivier Duroyon
Alcatel USA
15036 Conference Center Drive
Chantilly, VA, 20151
Phone: (703) 679 6415
Email: olivier.d.duroyon@usa.alcatel.com

Rudy Hoebeke
Alcatel Bell
Francis Wellesplein 1
2018 Antwerp, Belgium
Phone: (32) 3/240.84.39
Email: rudy.hoebeke@alcatel.be

Hans De Neve
Alcatel Bell
Francis Wellesplein 1
2018 Antwerp, Belgium
Phone: (32) 3/240.76.94
Email: hans.de_neve@alcatel.be

Dimitri Papadimitriou
Alcatel NSG-NA
Francis Wellesplein 1
2018 Antwerp, Belgium
Phone: (32) 3/240.84.91
Email: dimitri.papadimitriou@alcatel.be