

Open Database for Interconnected Traffic Engineered Multi-Layer Networks

F. Paolucci⁽¹⁾, F. Cugini⁽²⁾, G. Cecchetti⁽¹⁾, P. Castoldi⁽¹⁾,

(1) Scuola Superiore Sant'Anna, Pisa, Italy; (2) CNIT, Pisa, Italy;

fr.paolucci@sssup.it

Abstract: A database for interconnected TE networks is proposed and implemented to enable dedicated applications to perform complex multi-domain/layer operations. The database is successfully utilized for effective and highly-scalable maintenance applications in multi-layer networks.

OCIS codes: 060.0060, 060.4250

1. Introduction

The Path Computation Element (PCE) architecture [1] has been introduced to provide effective Traffic Engineering (TE) solutions by exploiting a dedicated network entity devoted to path computation process. Communications between path computation clients (PCC) and PCEs, realized through the PCE Protocol (PCEP), also enable inter-PCE communications offering an attractive way to perform TE-based path computations among cooperating PCEs having different domain of visibility. The multi-layer scenario is a typical example of different domains of visibility, mainly due to different technologies (e.g., optical or packet switching), different equipment vendors, different administrative management, and scalability issues. In multi-layer networks, correlation and exchange of TE information may be achieved by exploiting PCEP-based procedures. For example, a PCE controlling the Packet Switching Capability (PSC) domain may request a PCE dedicated to the Lambda Switching Capability (LSC) domain to expand loose Label Switched Paths (LSPs) in the optical domain. More recently, the Link State extensions to Border Gateway Protocol (BGP-LS) have been also introduced to provide the cooperating PCE with TE information, thus facilitating path computations. However, PCEP and BGP-LS protocols are typically designed to limit the amount of exchanged information (e.g., limited amount of segment expansions, exchange of address reachability information only), such that scalability issues are avoided on the remote elements. In addition, and most important, some pieces of information are often not dynamically retrieved by the control plane, such as the association between a card in the PSC router and the physically attached LSC transponder in a ROADM. For these reasons, several multi-layer operations, e.g. for maintenance and re-optimization, are still left to manual interventions, without relying on fully automatic mechanisms.

In this work, we take advantage of the impressive scalability performance of the latest generation of database technologies [2] to provide an effective and standard way to correlate TE information originated by different domains/layers. An OpenDatabase is implemented to enable dedicated management/control applications to perform complex TE operations. An automated maintenance application is implemented and evaluated in terms of scalability in a multi-layer PCE-based control plane testbed including Elastic Optical Network (EON) and Segment Routing MPLS domain.

2. Open Database Architecture and Use Case

In this section, the proposed solution, based on an OpenDatabase (OpenDB) for TE information correlation, is introduced. The use case of optical link maintenance in multi-layer networks is considered as a reference scenario. Single Operator scenario is assumed. The sample multi-layer core network composed of five sites is assumed, as shown in Fig. 1a. Each site includes a Router (equipped with PSC interfaces) and a ROADM (equipped with LSC interfaces). Each router interface is physically attached to a ROADM interface (i.e., transponder). In this example, lightpaths are established between adjacent ROADMs. Moreover, one additional lightpath is computed by the LSC PCE and established between ROADM A and C, having optical bypass in D. The obtained virtual topology at the PSC layer is shown in Fig. 1b. Note that a direct link is now introduced between Router A and C, enabled by lightpath A-D-C. Several MPLS LSPs are then computed by the PSC PCE and established, including an LSP from A to F passing through router C (LSP PSC 13 in Fig. 1b). Maintenance operation of an optical link is then considered. To avoid traffic disruption due to rerouting operations at the LSC layer, rerouting of involved PSC LSPs is typically performed in

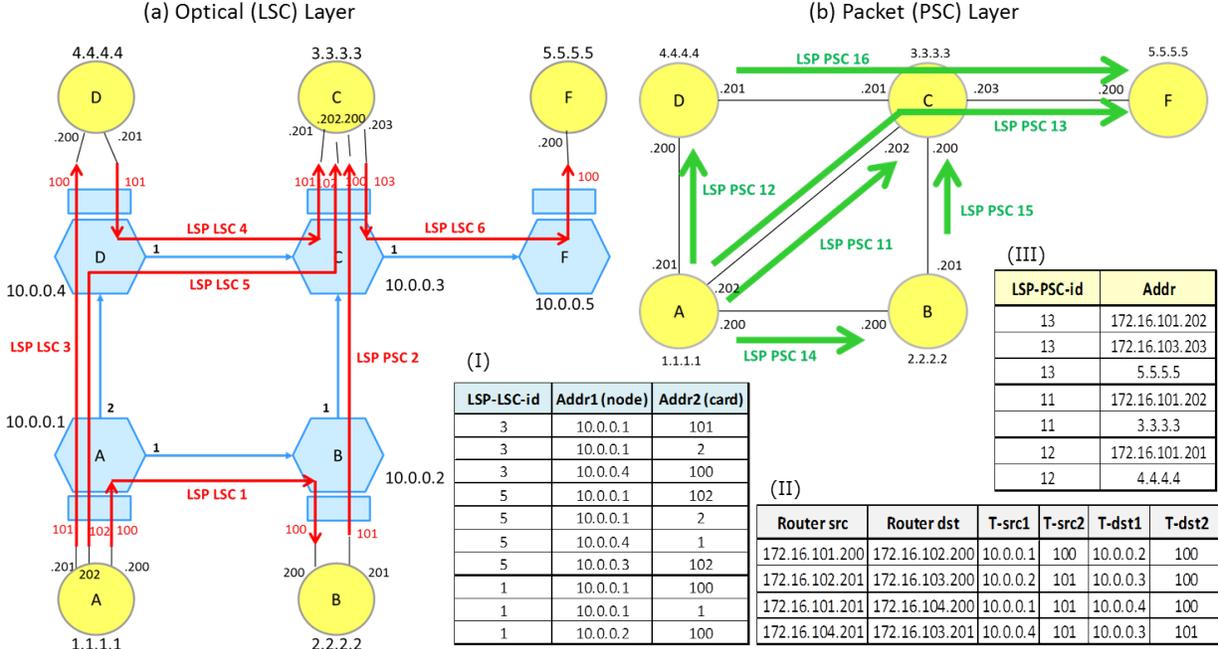


Fig. 1. (a) Optical (LSC) network: ROADMs, active LSC LSPs and inter-connected Routers; (b) Virtual network: PSC links and PSC LSPs; Table I: active LSC LSPs and traversed LSC interfaces; Table II: interconnection between LSC and PSC interfaces; Table III: active PSC LSPs and traversed PSC interfaces

advance at PSC layer, subject that adequate bandwidth resources are available in the rest of the PSC network. First, affected lightpaths have to be identified and associated to the related virtual links at the PSC level. Then, the PSC LSPs traversing these virtual links have to be rerouted. For example, in case of maintenance of link A-D, lightpaths A-D and A-C have to be identified, corresponding to virtual links A-D and A-C at the PSC level. Then, three LSPs have to be selected for rerouting: A-D, A-C and also A-F. However, the PSC PCE is not typically able to automatically perform the needed correlations between the actual optical links exploited by the PSC LSPs. Indeed, PCEP and BGP-LS may not be adequate, since the interconnection between router interface and ROADM transponder is not typically included in control plane advertisements and concerns on scalability issues are considered when providing each PCE with excessive amount of information (particularly in the case of active stateful conditions).

In this scenario, the proposed Open Database (OpenDB) is introduced. The OpenDB contains layer-specific information (e.g., active LSPs collected by each layer-specific PCE through PCEP PCReport messages) as well as inter-layer information, (e.g. static associations between the router interface and the ROADM transponder).

Standard YANG definitions should be considered within the OpenDB, enabling vendor-independent Applications (APPs) to perform queries and elaborations having access to selected parameters of the whole multi-layer network.

For the use case of link maintenance, a portion of the most relevant information included in the OpenDB is reported in Tables I, II and III of Fig. 1. As a first step, the APP devoted to maintenance procedures receives as input the optical link under expected maintenance. Then, it queries the OpenDB, collecting the association between lightpaths and traversed optical interfaces (Table I), between optical interfaces and attached router interfaces (Table II), and between PSC LSP and traversed router interfaces and virtual links. As outcome, the OpenDB provides the APP with (1) the list of links at the PSC level to be excluded by the rerouting process and (2) the list of involved PSC LSPs. For example, in the case of maintenance applied to link A-D, the APP queries the OpenDB as in the following. The A-D link described with its unnumbered interface, i.e. composed by the node IPv4 address (10.0.0.1) and the local identifier (value 2), is considered as input parameter. According to Table I, the OpenDB first identifies two LSC LSPs (i.e., LSC LSP Id 3 and 5) traversing A-D. Their LSC interfaces (i.e., 101 and 102 at 10.0.0.1) are used within Table II to retrieve the IP addresses of the attached interfaces at the packet layer (i.e., 172.16.101.201 and 172.16.101.202 respectively). Such IP addresses are used within Table III to retrieve the PSC LSP Identifiers utilizing these interfaces (i.e., ID 13 and 11) and the related LSP Explicit Route Objects (ERO). Note that all these correlations are performed without involving the PCEs. Only after this stage, the APP triggers a PCEP PCReq to request the PSC PCE to perform

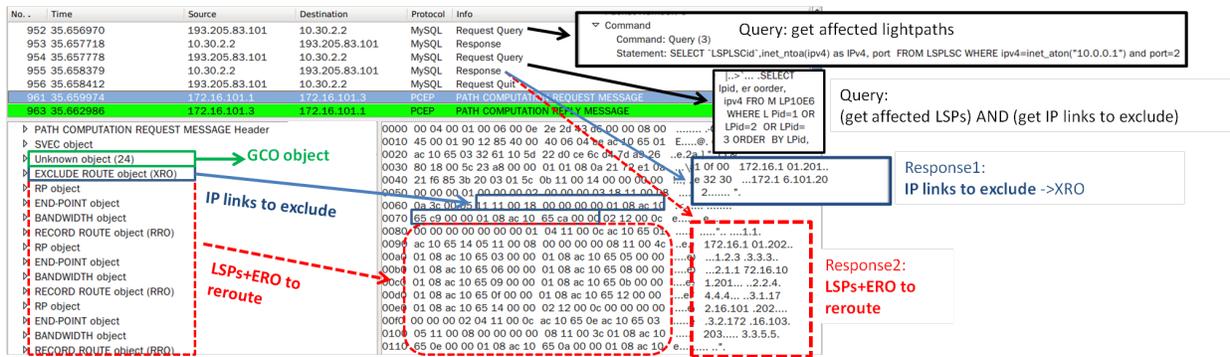


Fig. 2. Message exchange by the Maintenance APP with the OpenDB and PSC PCE

a Global Concurrent Optimization (GCO) of the identified involved LSPs (the three PSC LSP A-D, A-C and A-F), including the specific PSC links (A-D and A-C) within the eXclude Route Object (XRO). The PSC PCE then performs joint path computation and, in case of active stateful condition, it is also able to apply the computed changes, releasing the resources actually traversing the optical link A-D and redistributing them according to available resources.

3. OpenDB implementation and experimental demonstration

A network testbed implementing the key portions of the sample network shown in Fig. 1 is considered. The testbed includes the C++ LSC PCE implementation derived from [3] and supporting flex-grid EON and transponder extensions as in [4]. The considered C++ PSC PCE is derived from [5], also supporting Segment Routing extensions. The OpenDB has been configured as shown in Tables I-III on a MySQL platform (version 5.6) running on a remote FreeBSD server (AMD Opteron, 12-core CPU at 2.4GHz, RAM 16 GB). The maintenance APP has been implemented in C++. It includes a MySQL client for the communication to the OpenDB and a PCC for triggering the rerouting of selected PSC LSP. The database has been populated by the LSC and PSC PCEs upon any received PCReport message.

Fig. 2 shows the Wireshark capture of the full maintenance procedure of link A-D, as detailed in the previous section. Packets 952 to 956 report the queries and responses between the APP (IP addresses 172.16.101.1 and 193.205.83.101) and the OpenDB (IP 10.30.2.2). In particular, the first query (packets 952-953) retrieves the affected LSC LSPs, the second query (954-955) matches the related IP endpoint interface and the IP links to be excluded, also retrieving the PSC LSPs subject to rerouting and their ERO. Then, packets 961 and 963 report the subsequent PCEP session between the APP and the PSC PCE (IP 172.16.101.3), applying the rerouting of the retrieved PSC LSP A-D, A-C and A-F, subject to GCO and XRO excluding addresses 172.16.101.201 and 172.16.101.202 (i.e., router links A-D and A-C). The overall control plane procedure is completed in around 6ms, including around 1.4ms to perform the considered queries. To assess the scalability performance, 100k PSC LSPs over a 100nodes mesh network have been also emulated and inserted within Table III (reaching around 1 million entries). No scalability issues have been experienced: the time required to perform the last query for lightpath to LSP correlation always remained below 2ms.

4. Conclusions

An open database (OpenDB) for interconnected multi-layer TE networks is proposed, implemented and demonstrated for automated maintenance use case. TE parameters under different domains of visibility are properly correlated, and the overall maintenance is successfully completed at the control plane level in few milliseconds, with no scalability issues even with one million entries in the database.

Acknowledgment: Partially supported by the 5GEx project.

References

1. F. Paolucci, F. Cugini, A. Giorgetti, N. Sambo, and P. Castoldi, "A survey on the path computation element (PCE) architecture," *Communications Surveys Tutorials, IEEE*, vol. 15, no. 4, pp. 1819–1841, Fourth 2013.
2. L. Vaquero *et al.*, "Dynamically scaling applications in the cloud," *ACM SIGCOMM*, 2011.
3. L. Gifre *et al.*, "First experimental assessment of abno-driven in-operation flexgrid network re-optimization," *Lightwave Technology, Journal of*, vol. 33, no. 3, pp. 618–624, Feb 2015.
4. R. Martinez, R. Casellas, R. Vilalta, and R. Munoz, "Experimental assessment of gmpls/pce-controlled multi-flow optical transponders in flexgrid networks," in *OFC Conf.*, March 2015.
5. A. Sgambelluri *et al.*, "Experimental demonstration of segment routing," *J. of Lightwave Technology*, 2015.