The case for Open Lightpath Exchanges

Authors: Erik-Jan Bos, Lars Fischer, Tony Breach, Jerry Sobieski, René Buch
Date: January 25, 2012
Version: 1.0.1
Status: Final version

1. Executive Summary

Open Lightpath Exchanges perform the comparable function for lightpaths, that the Internet Exchanges -that emerged 15+ years ago- perform for the routed Internet. An Open Lightpath exchange is a policy-free interconnection point for lightpath-segments. These lightpath-segments each do have a policy attached to them, as they are scarce and costly resources most of the time. Connecting to an Open Lightpath Exchange is not necessarily free of charge, but is mostly based on a cost recovery scheme. Many Open Lightpath Exchanges nowadays exist around the world, and they are becoming a crucial building block for the lightpath capable parts of the NREN's hybrid networks, especially in emerging architectures such as iNDDI of Internet2 and other Software Defined Networking initiatives and in large and demanding users that create Open Network Environments, such as LHCONE.

2. Introduction

Roughly a dozen Open Lightpath Exchanges are up and running today, inside Europe, North America, South America, Africa and the Asia-Pacific region, each run by a (national) research and education network organization, or a related organization. Their number and sizes have grown steadily over the last ten years, initially loosely coordinated inside the Global Lambda Integrated Facility (GLIF). Now, ten years later, Open Lightpath Exchanges, also referred to as GOLEs[1], form a crucial building block in advanced networks for the emerging big-data e-Infrastructure, much in the way the Internet Exchanges have become a crucial building block for the routed Internet.

Open Lightpath Exchanges allow policy-free switching of end-to-end connections (lightpaths) delivered by multiple network service providers (connectors), using the facility for flexible hand-over when it comes to technology, operation and policy-neutral stitching. Hence:

- Open Lightpath Exchanges are technology-aware and inclusive of technologies, transparently acting as inter-connector or connector-translator between two or more lightpath-segments.
- Open Lightpath Exchange operation is lightweight; supporting traditional network management processes as well as emerging capabilities such as dynamic provisioning.
- Open Lightpath Exchanges are use-policy-free; cross-connects are established solely on the basis of bi-lateral agreement between the connectors requesting the cross-connection.
- Open Lightpath Exchanges are located in carrier-neutral housing facilities,

1 GOLE stands for GLIF Open Lightpath Exchange.

ensuring reasonable and non-discriminatory access into the facility.

An Open Lightpath Exchange is essentially a sub-IP peering fabric for organisations like (N)RENs and related organizations that want to stitch lightpath segments together. IP peering can of course be done over such a set of stitched lightpath segments, but in essence this is “higher layers” to the Open Lightpath Exchange itself.

Zooming into Europe, four open lightpath exchanges can be identified currently, i.e. from north to south:

- NorthernLight (responsible organization: NORDUnet)
- NetherLight (responsible organization: SURFnet)
- CzechLight (responsible organization: CESNET)
- CERNLight (responsible organization: CERN)

In Marseille, France, RENATER is working on creating MOLEN, the Marseille Optical Light Exchange Node.

The concept of an open lightpath exchange, just like with the concept of an Internet Exchange, scales to many locations around the world at places where three or more lightpath capable infrastructures come together [2]. However. it is expected that a small number of relatively large facilities on each continent, positioned in highly-connected locations, will emerge over time. Looking at Europe, a small number of new open lightpath exchanges is still required to serve the European region.

3. Hybrid networking vs. routed networking

The rationale for hybrid network, defined as IP routed services plus lightpath services, emerged roughly ten years ago when scientific applications started to foresee and announce the production of continuous data streams of bandwidths far exceeding what the NRENs of those days could offer or had planned. The concept of lightpath networking was born in 2001, as an economically and technically viable solution for the enormous data streams that would be feasible in the coming years. First discussed at the TERENA Networking Conference 2001 in Antalya, Turkey, in May 2001, this was more broadly discussed at the Global LambdaGrid Workshop in Amsterdam, in September 2001.

Ten years ago it was clear that given the price points of optical switches, Ethernet switches and IP routers (roughly 1 to 10 to 100 currency units for a 10 Gbit/s port) it would make great sense to provide connectivity as low in the transmission stack as possible [1]. Also separation of big continuous flows from small message driven Internet applications turned out to be essential, especially on long links with relatively high round trip times, as the prevailing Transmission Control Protocol stacks in use in the Internet suffer from properly handling the influence of the different types of traffic on each other within a shared infrastructure. Putting the big flows on their own lower layer segment of the infrastructure and allowing the use of aggressive (often UDP based) protocols turned out to be the solution for the larger eScience, eResearch and visualisation application based flows.

The price point differences have in the meantime eroded, but the rationale for hybrid networking is still present when we look at higher data speeds. For speeds at 10 Gbit/s a router port in e.g. a Juniper T1600 is still more expensive that such a port in an optical switch. For speeds beyond 10 Gbit/s, i.e. 40 Gbit/s, 100 Gbit/s, and beyond, it is clear that

router blades are expected to be extremely expensive when they hit the market. Given eScience is pushing the envelope with regards to high-volume data streams, we must continue to be prepared for these higher data rates; the hybrid model going forward will continue to be needed, supplying excellent routed connectivity for the modest data streams (e-mail, browsing, social networks but also video conference and modest cloud access) and lightpath-style networking for the demanding applications of the eScience and eResearch world, on separate lower layer paths on a common photonic transport infrastructure.

4. The role of Open Lightpath Exchanges in a connected world

Most of Europe has a one NREN per country landscape and has GÉANT to which all the NRENs are connected for inter-NREN connectivity. Moreover, GÉANT supplies worldwide connectivity for many of the European NRENs[2]. Open Lightpath Exchanges contribute to this connected world by offering "peering points" for lightpath segments of other world regions and for connecting to the non-NREN world, which often has services available for the NREN users, e.g. cloud services. For that reason, Amazon is now prominently connected to PacWave, the Open Lightpath Exchange at the West Coast of the USA. Hence, Open Lightpath Exchanges in Europe:

- Bring connection points for NRENs and related organizations of other world regions.
- Bring connection points for non-NREN type organizations delivering services to the (users of the) NRENs.
- Bring additional connection points for European NRENs, e.g. for adding resilience and richer connectivity, just like the Internet Exchanges do that for IP peerings.

Transit between Open Lightpath Exchanges on one continent could in theory delivered by anyone who can deliver high quality and high bandwidth point-to-point connectivity, suitable for transporting lightpath-segments. The European NRENs have commonly created the GÉANT network and GÉANT is a primary source for such connections for the NRENs. Cross-border fibers, an important architecture element for GÉANT and the GN3 Project, can be used as well of course for this purpose. See also the footnote on the next page on Internet2's strategic direction, positioning Open Lightpath Exchanges and transit between them as distinct products of the Internet2 service offering.

5. Policy handling at an Open Lightpath Exchange

An Open Lightpath Exchange is a policy-free lightpath-segment stitching fabric. The "policy-free" part of this definition might lead to confusion, however it only refers to the fact that the Open Lightpath Exchange itself does not add any policy, and will create a cross-connect between two connectors at the exchange when both are in agreement that such cross-connect should happen. As such the open lightpath exchange itself could also be called policy-neutral. It does not impose any policy over and above what already exist on the links coming into the exchange, and it assures that no traffic will flow to a connector at the Open Lightpath Exchange that is not approved by the connector.

From this it is clear that, regarding the risk that an NREN customer would directly connect to the open exchange as mentioned in the last paragraph of the previous

2 A number of European NRENs, besides their often rich connectivity to GÉANT, run international links themselves, e.g. to North America.

chapter, GÉANT and the NRENs are always in charge of who connects to GÉANT and the NREN. So, through policy mechanisms in GÉANT and in the NREN it can be discouraged that an institute for research and education would directly connect to an open lightpath exchange for its NREN-world connectivity. The same is true for an NREN that would only connect to an open lightpath exchange[3].

A policy free Open Lightpath Exchange has also requirements to its housing. The housing site owner should not prohibit parties from entering and connecting to the facility at fair and reasonable conditions.

Policy-free don't imply that the connector can connect to the Open Lightpath Exchange for free. The owner/operator of the Open Lightpath Exchange can of course declare a cost for each physical connection to its fabric. Most cost models in use at Open Lightpath Exchanges around the world today are just based on cost recovery, covering for the investment in equipment and cost of operation.

6. Conclusion

Open Lightpath Exchanges have emerged over the last 10 years on almost all continents, originating from NRENs and their constituencies. These OLEs form an important architecture element in the modern NREN networks, and way of building and using these networks. New developments on Software Defined Networking center for a great deal around Open Lightpath Exchanges, with a huge standardization effort inside the Open Grid Forum on the Network Services Interface (NSI), jointly supported by leading experts from Europe, North America and Asia. E.g. Internet2 is now showing the way to create a service on this with their newest proposals for Open Lightpath Exchanges across the USA build using architecture element such as OS3E and NDDI [4].

Further Reading

[1] Cees de Laat, Erik Radius, Steven Wallace, "The Rationale of the Current Optical Networking Initiatives", iGrid2002 special issue, Future Generation Computer Systems, volume 19 issue 6 (2003).
[2] Role of Open Exchanges in the evolution of global research and education networking; May 19, 2011; Bill St. Arnaud.
[3] Open Lightpath Exchanges, documentation from the research done by Prof. Cees de Laat (University of Amsterdam) and his SNE group. Available online at: http://ext.delaat.net/olex/index.html
[4] iNDDI Engagement Program (iNDDI-EP). More information available online at: http://inddi.wikispaces.com/

3 Internet2, through its NDDI and OS3E annoucements, has chosen the direction of creating a number of Open Lightpath Exchanges across the USA, and to offer Internet2's lower layer network as an efficient transit between exchanges. For more information, see: http://www.internet2.edu/network/ose/