The role of open exchanges in research networking

Project : GigaPort3
Project Year : 2011
Project Manager : Gerben van Malenstein
Author(s) : Stratix Consulting
Completion Date : 30-06-2011
Version : 1.0

Summary

For lightpaths to be connected across domains, points of interconnect are needed. The open lightpath exchange model has a number of advantages compared to a more traditional aggregator model: open lightpath exchanges stimulate innovation, enable fast and flexible lightpath set-up, and impose no restrictions on the users or the content. As the reach and the flexibility of open exchanges increases, more and more research will be able to benefit from this model. The emergence of dynamic multi-domain lightpaths, enabled through dynamic open lightpath exchanges, will create further opportunities for research.

Colophon

Programme line	: GigaPort3
Part	: Work Package NetherLight and international connectivity
Activity	: INT-11-5
Deliverable	: Paper: The role of Open Exchanges in research networking
Access rights	: Public
External party	: Stratix Consulting

This project was made possible by the support of SURF, the collaborative organisation for higher education institutes and research institutes aimed at breakthrough innovations in ICT. More information on SURF is available on the website www.surf.nl.

6 Matters one should know about the role of Open Exchanges in research networking

Scenario	The emergence of lightpaths across multiple domains has also led to a need for interconnections between these lightpaths. One model to provide these interconnections is the open exchange model.
What is it?	An Open Exchange is a point where multiple connectors meet to transfer data amongst them on dedicated bandwidth.
Whom is it for?	Those who want to facilitate large data flows through dedicated bandwidth in a future-proof way.
How does it work?	The open part of 'Open Exchange' means that no restrictions are posed on whom connects to whom within the Open Exchange. This way, members connecting to an Open Exchange will not be constrained by any rules.
What can one do with it?	Exchange large data transfers with more than two entities in a future-proof way.
More information	http://www.netherlight.net

Index

1 Introduction

Research is increasingly becoming a global effort, based on international co-operation and resource sharing. For optimal use, more and more components of the science infrastructure will be geographically dispersed, and used in different configurations over time. Large instruments, storage and computing facilities are becoming generic resources that can be combined to provide services through the network infrastructure to an increasingly diverse research community, while the huge data-sets that are generated in all types of research need to be available to the research community at large.

These trends have led to the need for fast, static or semi-static point-to-point connections between endpoints, usually known as lightpaths. In recent years, an increasing number of research institutions have started to use these lightpaths, not only within a single network domain but also across domains.

The emergence of lightpaths across multiple domains has also led to a need for interconnections between these lightpaths. One model to provide these interconnections is the open exchange model. Open exchanges allow networks to interconnect at neutral meeting points. These meeting points are "policy free", in that the exchange does not impose any restrictions on the type or amount of traffic exchanged by the participants.

This paper explores the emergence of the open exchange model, and analyses how this model helps satisfy the needs of the research community.

2. Background: the Lambdagrid

2.1 The changing research environment

The global research environment is changing rapidly. There is an increased role for ICT in different scientific disciplines, which is not limited to the beta sciences but extends to all disciplines.
Research generates increasing amounts of data which need to be transported, processed, stored and retrieved. At the same time, research is already a global, collaborative effort, combining the efforts of researchers situated in different places, and it is becoming ever more so.
Major instruments are also becoming more distributed (as in the case of sensor networks) or centralised (as in the case of high energy physics), in both cases leading to a geographical separation between researchers, processing centers, and instruments. All of these trends create a need for data connections which are faster and more reliable than ever before, while remaining affordable within the limited budgets available.

2.2 The case for lightpaths

National Research and Education Networks (NRENs) have traditionally[1] provided ever larger IP connections to the institutions for research and education in their respective countries. SURFnet, for instance, currently provides uncongested 1 and 10 Gigabit/s IP connections to its connected institutions.
However, the developments mentioned above have resulted in a need for a different type of service: in addition to the existing "best effort" IP service, there is a requirement for dedicated, unconstrained capacity between specific endpoints, with minimal latency and jitter, often at multiple Gigabits per second between a limited number of endpoints. This need has led to the development of lightpaths: fixed paths through the network, with dedicated capacity, low latency and practically no jitter.
Technically, lightpaths can be created in several ways; in most cases a Lightpath is either a wavelength (lambda) through an optical network or a transport path within such a lambda[2]. Current technology allows for lambdas transporting 10, 40 or 100 Gigabits per second; at this time most lambdas have been implemented as 10 Gbit/s optical paths but there is a growing tendency towards 40 or 100 Gbit/s. Where lightpaths are created as transport paths within a lambda, they can have different capacities depending on the technology used, but 150 Mbit/s and 1 Gbit/s are fairly common.
As NRENs started offering lightpaths to their users, more applications were found for these lightpaths. For instance, institutions with branch offices discovered that a transparent service between their routers made it easier to link branch offices into a single network, projects creating computing grids were able to link multiple data centres into a single grid, and distributed scientific instruments were linked to create virtual large scale instruments.
Lightpaths are becoming an important part of the larger infrastructure of e-Science. The e-Science infrastructure's goal is to be the ICT catalyst that facilitates collaboration between different scientific domains; this requires not only networking facilities but also high performance computing, storage and

---

1 At least since the early nineteen-nineties
2 For instance, through an OTN container, an SDH container or an Ethernet pseudowire over MPLS-TP.

resource management components which together provide a set of generic e-Science services, available for the specific demands of different types of research. This principle is illustrated in figure 1.

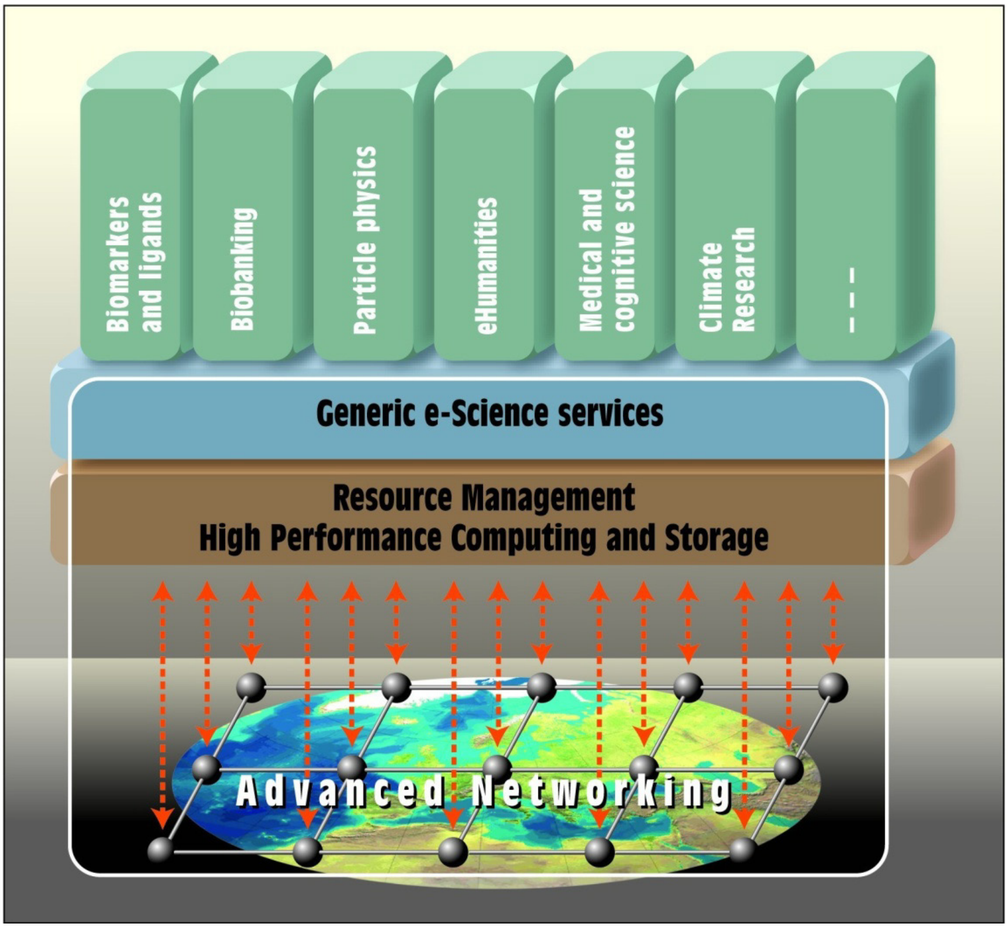

Figure 1: An infrastructure for e-Science (source: Herzberger, Sloot et al.)

Within this e-Science infrastructure, advanced networking provides the "glue" that links all the other components. For some parts of the infrastructure, the "best effort" IP service may be sufficient, while others require the predictability, transparency, low latency and high capacity provided by lightpaths.

2.3 Lightpaths across domains

As mentioned above, scientific research increasingly requires world-wide collaboration. Scientific research is thus not confined to single country or to the reach of any single network, and the e-Science infrastructure likewise needs to become global in nature.

As a result, Lightpath services are necessary which not only connect end-points served by the same network, but also end-points in different parts of the world and served by different networks. Therefore, lightpaths have to be provided across multiple networks or domains. This is illustrated in figure 2.

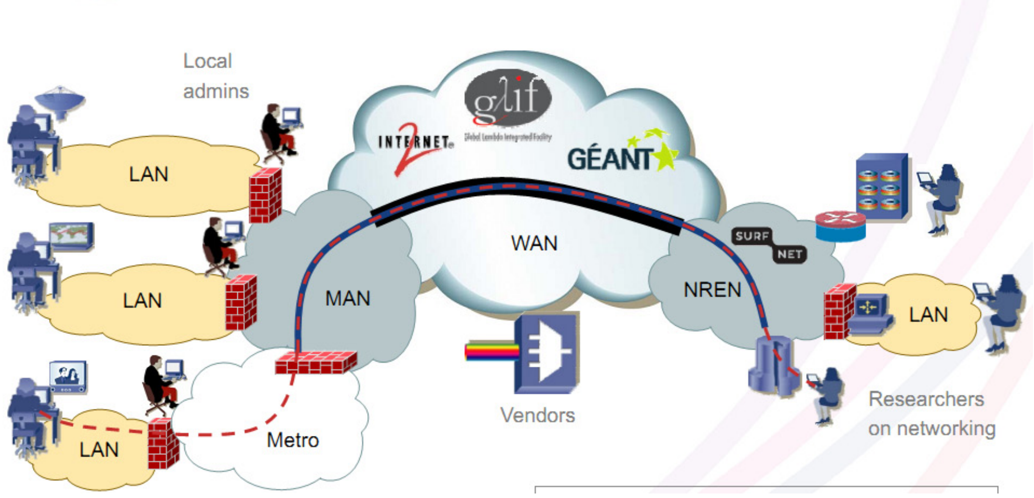

Figure 2: Multi-domain lightpaths *(source: Terena End-to-End (E2E) Provisioning Workshop)*

The need for lightpaths across many domains has led to an evolving grid of connections based on interconnected lambdas, often referred to as the *Lambdagrid*.

3. Interconnections in multi-domain networking

Lambda connections between networks, and end-to-end lightpaths across those connections, can be arranged on a bilateral basis between networks. However, this model is not very scalable as it would required each network to make the necessary organisational, technical, and financial arrangements with every other network to which it needs a connection. There is therefore a need for a more scalable model.

From the history of global telephony and data networks, it is clear that there are two basic models to avoid the large number of bilateral arrangements mentioned above: through open exchange points or through aggregator networks.

3.1 Open exchange model

In an open exchange model, each network is connected to a limited number of exchanges. These exchanges are points of interconnect for networks, allowing unrestricted connections between them. An exchange is, in principle, a single point to which links from multiple parties can connect. These parties can then use the exchange to provide connectivity as well as other services to anyone connected to the exchange.

Open exchanges offer "policy free" cross-connects among the ports of an exchange. "Policy free" means that there are no conditions imposed by the Exchange Point governing who can connect to whom or what kind of traffic is carried across the connection.

The exchanges, in turn, are connected to others networks or to other exchanges; the connections between may be operated by the owner[3] of the exchanges, but other users can also add connections. A model based on this type of exchanges is shown in figure 3.

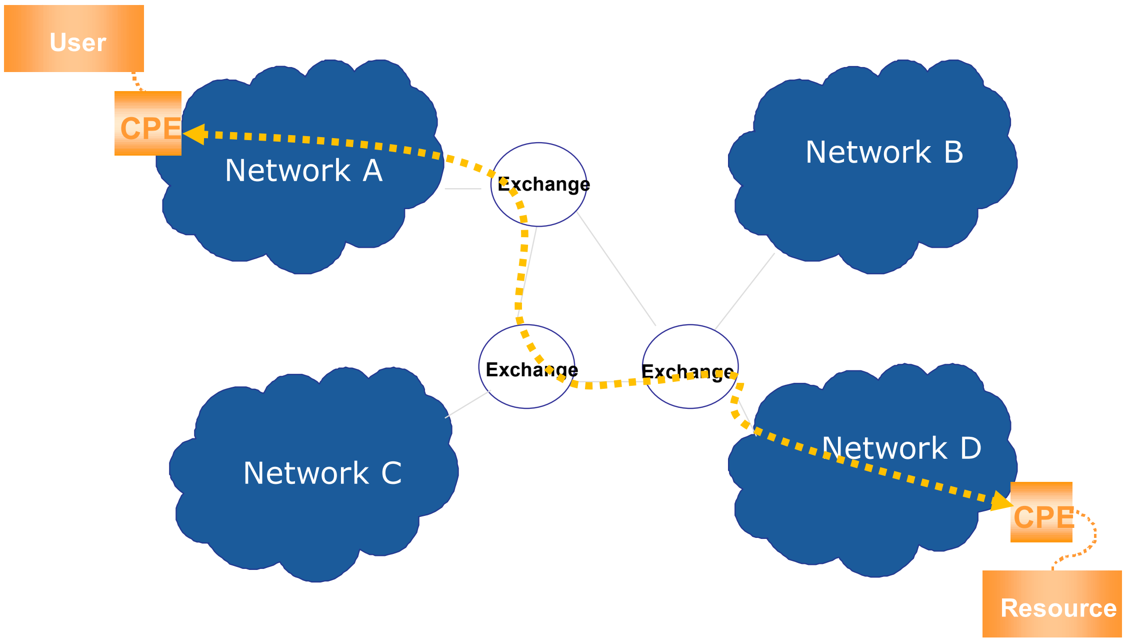

Figure 3: Multi-domain network using exchanges

[3] "Owner" in this context refers to the economic owner of a resource. This entity may be separate from the legal owner, for instance through a lease arrangement.

The open exchange approach has become standard for IP interconnection (the public internet). A well-known example of such an exchange is the Amsterdam Internet Exchange (AMS-IX). Due to the emergence of open internet exchanges, internet service providers are able to peer with each other without being restricted by the larger transit operators, as was the case in the remote past when most internet exchanges were part of the large carriers.
With the increasing demand for multi-domain lightpaths, in particular for research, an open interconnection model is needed, similar to the open internet exchanges, but on a lower layer in the "stack".
An open exchange model can create a number of advantages:

- As the exchanges are independent of the connections between exchanges, this model allows for flexibility and competition. If there are multiple types of connections available between the exchanges, from different parties and using different technologies, different approaches can be tested and compared. This creates an incentive for all the involved parties to experiment in order to find the most effective approach.

- As the exchanges have a clearly defined, limited role, the size and scope of the exchanges is easier to manage compared to an aggregator network, even as the number of connections grows. An open exchange model does not require every exchange to be linked directly to every other exchange, as long as there is always a possible path, and sufficient capacity, to every other exchange.

- Open exchanges, by allowing links from all parties to connect, enable these parties to share resources with each other. For example, a network or a research institution owning a lambda on one route can swap capacity with another institution on another route.

- Users requiring higher availability can decide to implement multiple routes between their endpoints, through different exchanges. This provides a choice of service levels for the user.

- The open exchange model allows intermediary parties (brokers) to combine different elements, such as connectivity, processing and storage, into a complete end-user service. The user can choose any broker, and anyone can become a broker.

- By definition, an open exchange has a virtually unlimited capacity (and can be easily expanded if this is no longer the case). The links between exchanges may become congested, but as there are multiple parties providing these links, there are sufficient opportunities for the user to acquire capacity.

3.2 Aggregator model

In an aggregator model, small networks make arrangements with one or more aggregator networks covering a larger geographical area, which in turn have arrangements with a limited number of other aggregator networks. Such a model is illustrated in figure 4.

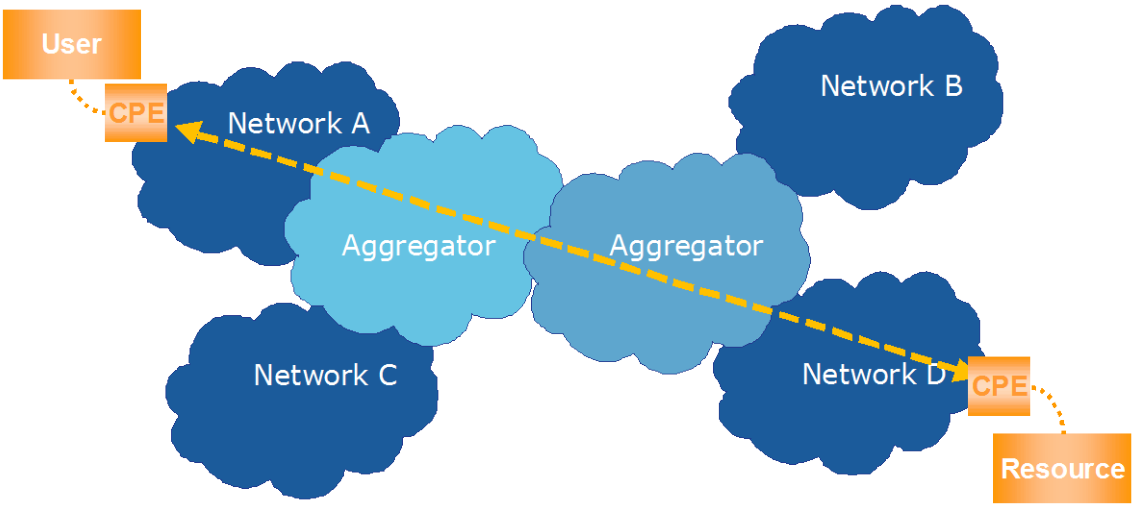

Figure 4: Aggregator model

This model has a number of advantages:

- This model can, when managed properly, provide global connectivity with a minimum of hassle; the aggregator arranges all needed organisational and technical issues for the user.

- This model is characterised by having large scale aggregators who in turn connect to multiple other networks. Therefore, the aggregator can achieve more economy of scale, in particular when the individual traffic flows are small, and therefore offer cheaper services.

- There is operational ease for users, because all traffic can be handled by the chosen aggregator.

However, the aggregator model has also a number of disadvantages:

- Once the owner of a small network (e.g. a campus network) has invested in a physical connection to one aggregator network, the investment forms a barrier for it to switch to another aggregator network. Therefore, there is little incentive for the operator of the aggregator networks to innovate, as there is limited competition between different providers or between different approaches.

- As the number of lightpaths grows, the aggregator networks have to become bigger and bigger. This makes it difficult for these networks to maintain flexibility. This is less the case for an open exchange model, as the function of each exchange and each link remains simple.

- The available capacity is completely dependent on the capacity of the aggregator networks. If this is insufficient, and can not be upgraded in time, the aggregators are forced to define admission policies, which may block some types of use. When multiple aggregators are involved, with conflicting policies, many types of use may be blocked, creating a barrier to innovation. In the open exchange model, the constraining factor is formed by the links, but as any party can bring a link into the exchange, any capacity constraints can be resolved either by the user or by any party interested in connecting the exchanges.

- If the networks in the lightpath all have their own "acceptable use" policies, the use of the lightpath will be restricted to the lowest common denominator of these policies. This will exclude many possible applications, even when there is enough capacity. In an open exchange model, the use is determined only by the policies of the links involved; as long as there are links available which permit a certain use, the lightpath can be routed through those links.

3.3 Combining both approaches

The combination of open exchanges and links between them is particularly suitable for large flows of traffic between specific endpoints. For smaller amounts, there is still a need for "aggregator" networks who will combine traffic from a large number of sources to achieve economies of scale. Access to such aggregators can either be through proprietary, private connections or through open exchanges.

If aggregator networks are connected to open exchanges, the users get the best of both approaches: flexible access to links to connect to other exchanges for large flows, and transit services for any destination in the world for smaller traffic flows. Such a hybrid situation is shown in figure 5: Network A is connected to an open exchange, which provides connectivity to Network B through a link to another open exchange, but also to an aggregator network which in turn connects directly to Network C and, through a third open exchange, to Network D.

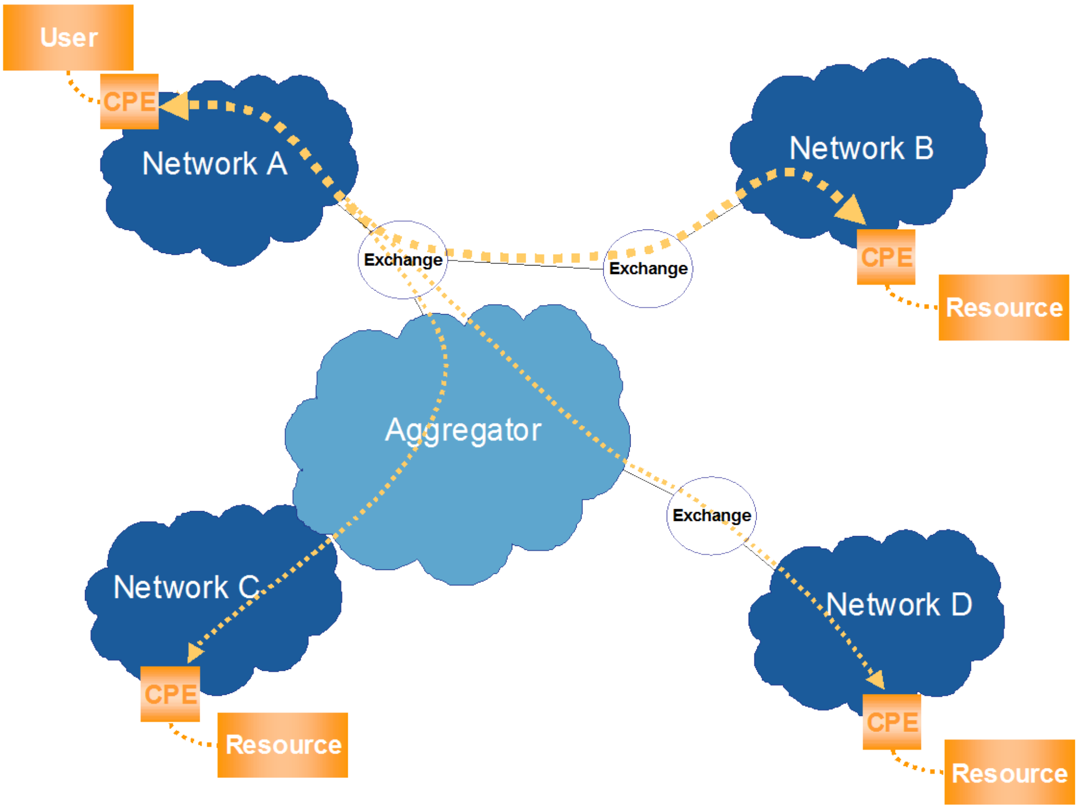

Figure 5: Combining the open exchange and aggregator approaches

4. The open exchange model in the NREN environment

Given the large data flows created in modern research, the open exchange model is particularly useful in the research environment. It is therefore not surprising that this model has been pioneered by NRENs, first for IP (in the nineteen-nineties) and more recently for lightpaths.

4.1 The GLIF

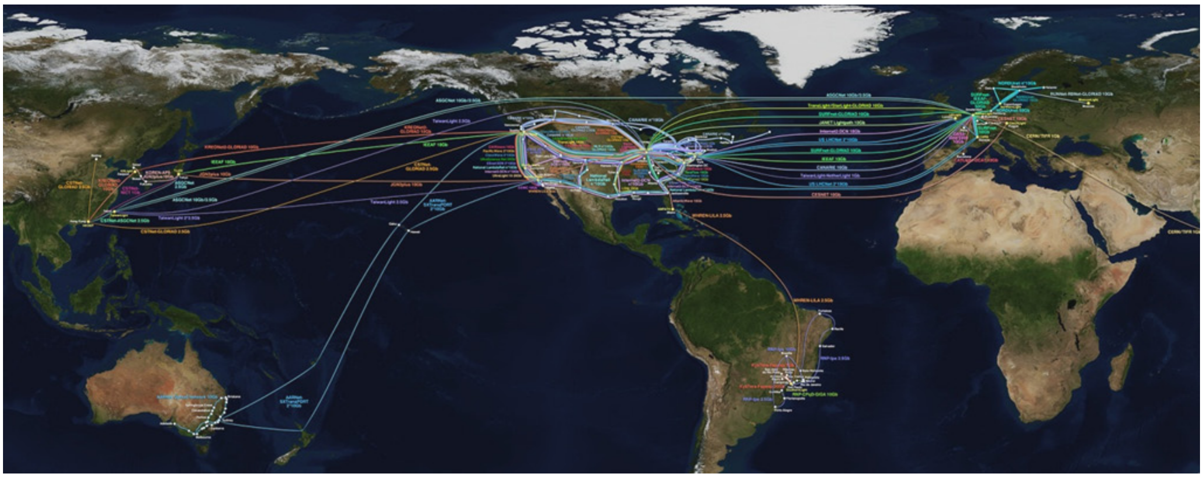

Figure 6: Open exchanges and links in the GLIF (2008, source: GLIF[4])

The vision of open, neutral peering facilities for lightpaths is shared by a large number of NRENs and related institutions around the world. As a result, there is now a global set of interconnected Open Lightpath Exchanges, including NetherLight, CzechLight en CERNLight in Europe, co-ordinated in the Global Lambda Integrated Facility[5] (GLIF). Within the GLIF initiative, there are now 19 GOLEs (GLIF Open Lightpath Exchanges).
An Open Lightpath Exchange is a peering point, separate from the network domains to which it is connected. NetherLight, for instance, is owned and operated by SURFnet but it is not part of the SURFnet network. Connections can be set up through the NRENs, traversing one or more Open Lightpath Exchanges to create the end-to-end connection, as illustrated in Figure 7

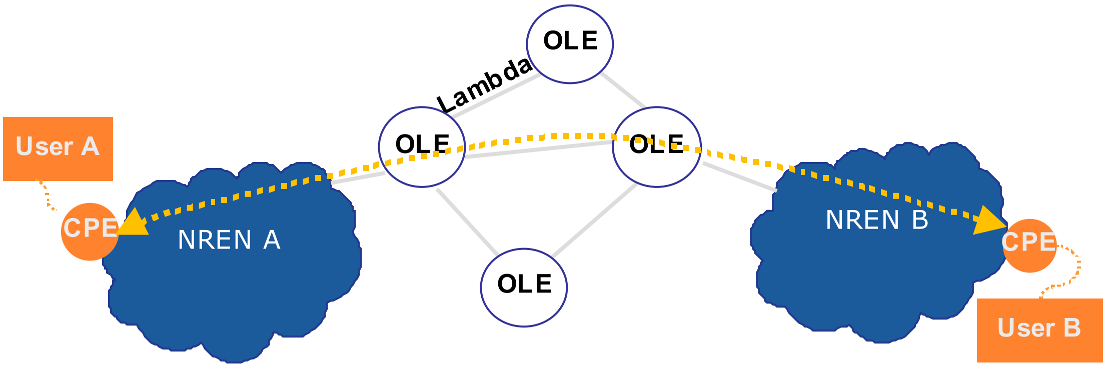

Figure 7: End-to-end connection across Open Lightpath Exchanges (OLEs)

Operators of Open Lightpath Exchanges around the world are currently working within the GLIF to create a set of end-to-end management mechanisms, in order

[4] http://www.glif.is/publications/maps
[5] Refer to www.glif.is

to facilitate easier lightpath set-up and operations. The lightpath management mechanisms currently being defined are, like the Open Lightpath Exchange structure itself, based on peer-to-peer co-operation rather than any form of central management.
Once connected through Open Lightpath Exchanges, users can easily set up temporary or permanent interconnections between them. For researchers, who often need to set up distributed infrastructures for experiments on a temporary basis, this creates a level of flexibility which is hard to achieve using more traditional models such as aggregator networks.

4.2 Organisational aspects

While open exchanges are not part of the NRENs' networks, in practice they may still be managed by the NRENs. However, the NRENs manage these exchanges as separate entities in their organisations; this ensures openness and transparency. In the long run, they may well decide to split off the open exchange as a separate organisation, as has happened with many of the internet exchanges in the past.
Besides guaranteeing the neutrality of the exchanges, the independence from the NRENs' networks also allows these exchanges to connect commercial parties as well as NRENs. Everyone is able to participate in an open exchange point, whether it is to provide links or services, to use these links or services, or some combination of these.
In an open exchange model, ownership of the exchange and of the links connecting the exchanges are also separated[6]. While some of the links connected to an exchange may be owned by the same entity owning the exchange, there should be no difference in treatment between those links and links owned by other parties.

4.3 Technical aspects

Open exchanges are connected (in most cases) through one or more lambdas, each at 10, 40 or 100Gb/s. Lightpaths across the exchange can share lambdas, so lightpaths can be smaller; usually lightpaths can have a bandwidth ranging from 100 Mb/s to 10 Gb/s. However, a lightpath can also consist of an entire lambda at 10 to 100 Gb/s.
The lightpaths that need to be connected can use different technology solutions[7]. Therefore, the exchanges must be able to connect different types of connections. Most lightpaths are currently SDH containers (using VCAT and GFP), so an SDH crossconnect is sufficient, but there are other technology solutions used by lightpaths, such as Carrier Ethernet pseudowires. An Open Lightpath Exchange may need to provide conversions services between various technologies.
Figure 8 shows a few examples of conversions within an actual lightpath between France and Japan.

6 For a detailed discussion different aspects of ownership in open exchanges, refer to "A Terminology for Control Models at Optical Exchanges", Dijkstra et al, july 2007, in LCNS Volume 4543 http://ext.delaat.net/pubs/2007-c-9.pdf

7 For an overview, refer to "Optical Exchanges", Dijkstra and de Laat, GRIDNETS 2004 (http://www.broadnets.org/2004/workshop-papers/Gridnets/DijkstraF.pdf)

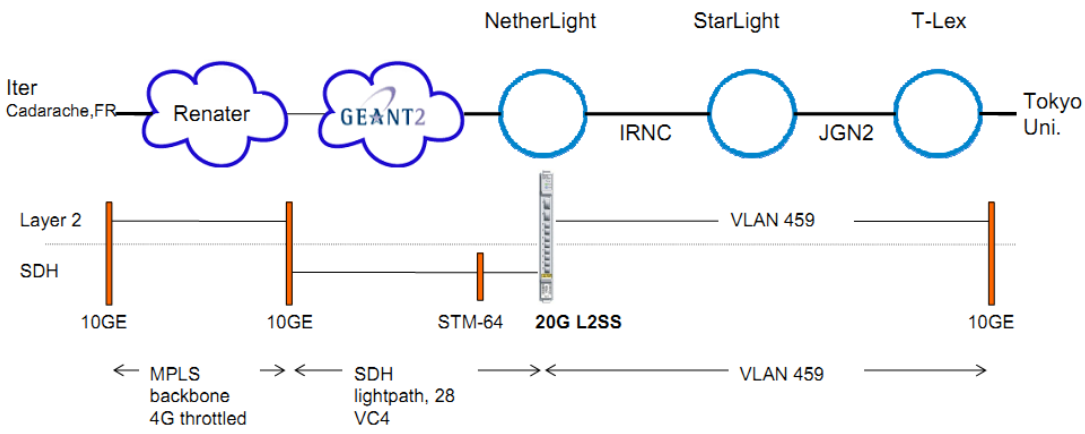

Figure 8: Conversions within a lightpath (source: presentation W. Huisman, TERENA 2009)

An open exchange can be physically distributed. However, in order to maintain the advantages of an open exchange, a distributed exchange would need an open policy and virtually unlimited bandwidth between the ports of the exchange. This is only feasible within a small geographical area, with affordable fiber connections between the nodes. Otherwise, it is actually a network, with its own constraints, rather than an exchange point.
An example of a distributed open exchange is the AMS-IX, which has ports in eight locations spread across Amsterdam, without imposing any restrictions on traffic between these locations. This only works because the scope is limited to a single city, with plentiful dark fiber available between sites.

4.4 Financial aspects

The costs of an open exchange are determined, for the largest part, by the number and capacity of the ports. There is therefore no need for any form of usage based charging; in a transparant financial structure, each party connecting to a port simply pays for the use of that port. This reduces the barriers to the actual use of the exchange, as there is no additional cost for additional use of an existing port.
The links between the exchanges are arranged and financed separately; a party needing a connection across multiple exchanges may be able to use capacity made available by any network owner that participate in the open exchange model. If there is no suitable link available, any party can acquire its own links. Within the context of the GLIF, there is a large number of links already made available for research and education purposes. For instance, between NetherLight (Amsterdam) and MANLAN (New York) there are currently links provided by IRNC, Canarie, SURFnet, and USLHCNet, Users outside the research and education field can still use the open exchanges, but may need to acquire separate links.

5. Use cases

The effectiveness and the flexibility of the Open Lightpath Exchange model is best demonstrated through the current use cases. This chapter provides a few examples.

5.1 Large Hardon Collider Open Network Environment (LHCONE)

The Large Hadron Collider (LHC) is the large particle accelerator at CERN near Geneva, which started initial operations in 2008 and is now fully operational. The huge amount of data that the LHC produces is transmitted to research institutes world-wide for distributed processing and analysis, through a set of lightpaths (most of them 10 Gb/s) constituting the LHC OPN (Optical Private Network).
Once processed at these major research institutes, known as Tier-1 sites, the data needs to be made available to a much larger group of smaller research institutes, the Tier-2 and Tier-3 sites. While the LHC OPN is a static network, the Tier-2 and Tier-3 sites need to be able to set up connections to multiple Tier-1 sites, as well as Tier-2 and Tier-3 sites. Therefore, a more flexible set-up was needed.
The LHC Open Network Environment (LHCONE) was designed to provide a flexible framework for the interconnection of Tier-2 and Tier-3 sites with each other and with the Tier-1 sites. The model is based on open exchanges and lightpaths, combined with aggregation networks for those sites not connected to an open exchange.
At this time (April 2011), LHCONE includes three open exchanges forming the initial core deployment: CERNLight, NetherLight, and Starlight. This core deployment already connects a number of Tier-2 centres on different continents.

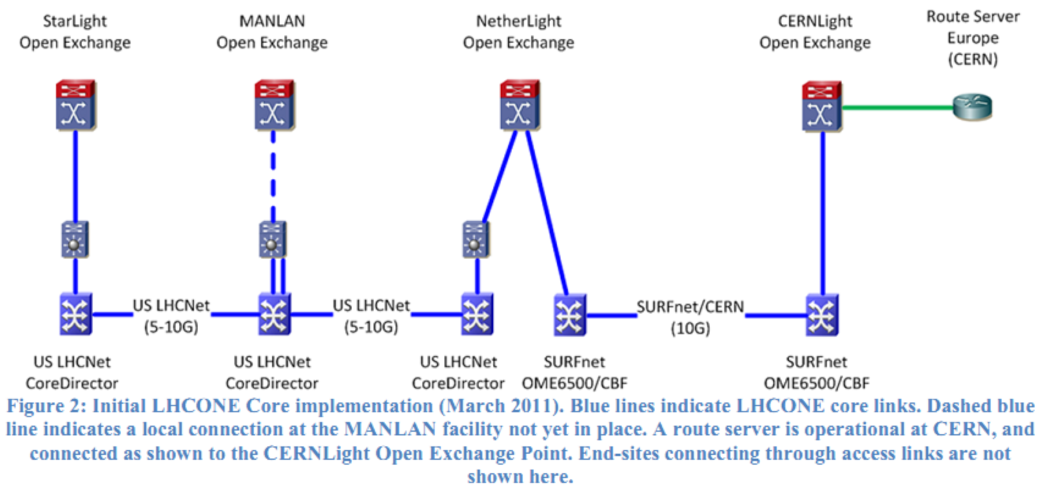

Figure 9: The initial LHCONE core

5.2 FEI, Eindhoven <-> USA

The Technology University of Eindhoven owns a high end FEI electron microscope, which it shares with users in other locations. FEI, the producer of the electron microscope has developed technology (both software and hardware) to make sharing easy and possible. Through lightpaths between the electron microscope in Eindhoven and researchers in the US, it is possible for a researcher

to operate the microscope from the US. This has been shown to work in demonstrations, and will soon be implemented as a production service. The lightpaths, connecting US researchers to the instrument in Eindhoven, are easily set up through the open exchanges in the Netherlands and the US, using link capacity already available for research purposes.
By sharing research instruments between universities, it is possible to share the investment costs of high end instruments and to collaborate in research that uses these instruments.

5.3 JIVE / e-VLBI, Radio telescopes

Very Long Baseline Interferometry (VLBI) is a technique that enables astronomers to observe objects in space at a very high resolution using radio frequencies. A data processor – known as the correlator – combines the signals received by an array of radio telescopes spread across the world. The telescopes produce large data streams during observation; at this time most of the radio telescopes available produce either 512 Mb/s or 1024 Mb/s, with some telescopes already able to produce 4 Gb/s. As this data can not be compressed before processing, the only practical way to bring this data together used to be the physical transport of large packs of hard disks.

Figure 10: The NEXPReS project links radio telescopes from around the world (source: Paul Boven, JIVE[8])

JIVE, the European consortium that manages the European VLBI Network Data Processor, is creating a network of lightpaths linking telescopes to the JIVE correlator (e-VLBI). At this time, most of the telescopes within Europe and several outside of Europe have been connected to the JIVE correlator through open exchanges. Other countries outside Europe and their local NRENs are currently working to connect their telescopes.
In the NEXPReS project, this concept is expanded to include more radio telescopes and dynamic bandwidth assignment.

---

[8] From http://www2.surfnet.nl/bijeenkomsten/rd2010/sheets/boven.pdf

6. Future developments in open exchanges

6.1 Next step: dynamic lightpaths across the world

As the global set of Open Lightpath Exchanges expands, it is becoming easier for NRENs to set up end-to-end lightpaths. While lightpaths will never entirely replace IP services, a large number of applications, for which high point-to-point bandwidth and low latency are important, will bypass the public IP cloud altogether.

A major next step in the improvement of lightpath networking is the introduction of dynamic lightpaths. Already, several NRENs have implemented dynamic lightpaths as a mechanism to provide point-to-point lightpaths on demand, with the end-user directly in control. Now, this mechanism needs to be expanded to allow multi-domain dynamic lightpaths across the world.

To this end, the global NREN community is working on collaborative approaches to dynamic end-to-end lightpath set-up. In this vision, the control planes of user institutions, NRENs and Open Lightpath Exchanges are linked on a peer-to-peer basis, providing the end-user with full control over the lightpath while allowing NRENs to manage their resources as needed. This mechanism is illustrated in figure 11.

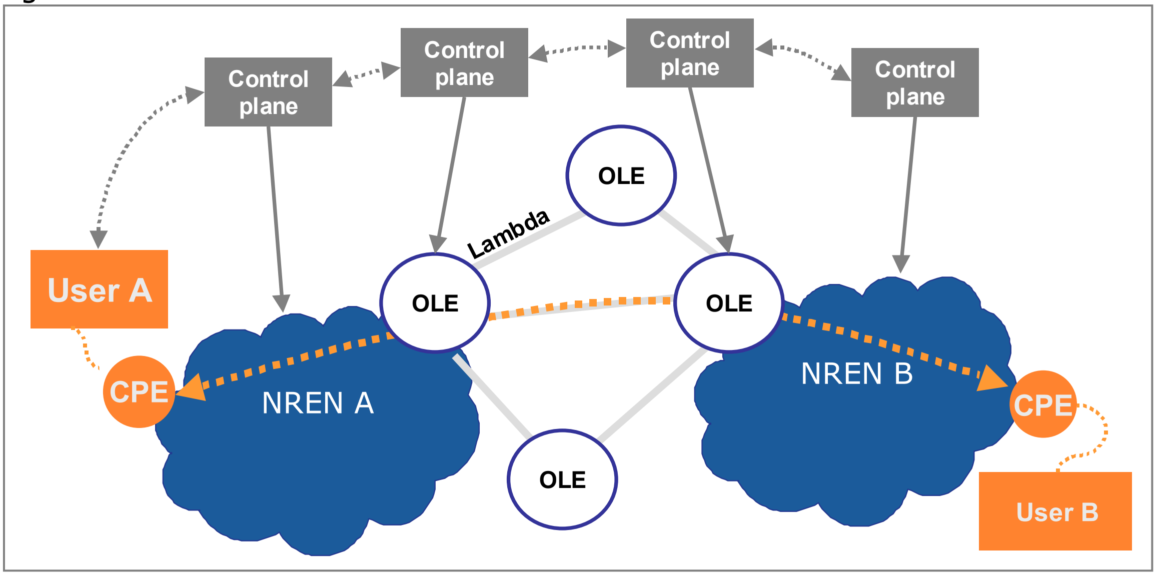

Figure 11: End-to-end dynamic management of lightpaths

Several implementations of dynamic lightpaths have already been linked to create multi-domain services; the community is now working on establishing a generic network services interface, so that different solutions can interwork with each other. This effort is being co-ordinated in the Open Grid Forum NSI-WG[9] (Network Service Interface Working Group).

Dynamic lightpaths are very relevant to projects such as NEXPReS, which require high bandwidth connectivity between a changing set of instruments and a central computing and storage system. Moreover, dynamic lightpaths also allow these central computing and storage platforms to become distributed and to be shared across multiple research institutions.

---

[9] Refer to http://www.ogf.org/gf/group_info/view.php?group=nsi-wg

6.2 More conversions at the Open Lightpath Exchange

Open Lightpath Exchanges already provide conversion between transport mechanisms. However, with Ethernet gaining importance as a transport technology, some new issues arise with respect to converting different types of Ethernet headers (such as 802.1Q, 802.1ad, and 802.1ah), the re-tagging of customer and provider VLAN identifiers[10], and the conversion between MPLS and Ethernet transport. Emerging standards such as the MEF E-NNI specifications might provide some guidance.

With the introduction of standard protocols to schedule multi-domain dynamic lightpaths (NSI), the conversion capabilities of each Open Lightpath Exchange will have to be communicated as part of the topology, so that the necessary conversions can be defined as part of the path finding algorithm.

6.3 Moving from experimental to professional

While the current Open Lightpath Exchanges were originally set up to enable experiments with lightpaths, the majority of traffic is now real production traffic. This requires a professional approach to the management of the exchange. Therefore, the open exchanges are currently improving their operational procedures in order to enable production quality services.

As the amount of production traffic increases, the governance and management of the open exchanges will need to become more professional to meet the needs of the users. The issues that need to addressed include the day-to-day management of the open exchange (with a focus on configuration management and fault management), the financial management, and a governance structure which ensures that open exchanges remain open.

6.4 Increasing the reach

Currently, there are 19 GLIF Open Lightpath Exchanges (GOLEs) around the world. While it may not be effective to have hundreds of GOLEs, the number of GOLEs can be expected to increase until every research institution around the world can have access, at reasonable costs, to at least one GOLE. Some major research facilities may even start up their own GOLE, as has been the case with CERN in setting up the CERNLight exchange. Others may want to arrange for a lambda or dark fiber to the nearest GOLE; for instance, the JIVE facility in Dwingeloo uses a DWDM system to connect to NetherLight, providing it with near unlimited capacity to any research institute or radio telescope which can arrange connectivity to NetherLight. As the number of GOLEs and the number of links between them increases, such connectivity becomes easier to arrange[11].

The open exchange idea is also gaining interest in the commercial environment. While carrier neutral, open internet exchanges have been the norm for over ten years, the increasing role of Ethernet as a transport technology has also led to a growth in commercial open Ethernet exchanges such as Equinix, Neutral Tandem and CENX. Aggregator-like services are also developing commercially.

While the use of Open Lightpath Exchanges has so far largely been confined to the research community, in most cases there is no reason why commercial organisations should not be able to use these exchanges. Obviously, this does not apply to the links between exchanges, which have been made available

---

[10] For details, refer to http://www.terena.org/activities/e2e/ws2/slides1/2_SURFnet_wouter.pdf

[11] Obviously, this model in itself does not solve the issue that most radiotelescopes have limited data connectivity due to their remote location. However, the model does ensure that once a radiotelescope is connected to the nearest well-connected GOLE, it can be accessed without restrictions.

specifically for research, but a commercial organisation could acquire its own links into an Open Lightpath Exchange and use this to connect to institutes for research and education. In an environment where these institutions are increasingly outsourcing some of their IT functions to commercial organisations, this may be an effective way to achieve the connectivity required.

7. Conclusions

With the increased use of ICT infrastructure by the research community, and the widespread international collaboration in research, there is a growing need for global, high bandwidth connectivity. Lightpaths, including multi-domain lightpaths, play a major role in satisfying this need.

For lightpaths to be connected across domains, points of interconnect are needed. The open lightpath exchange model has a number of advantages compared to a more traditional aggregator model: open lightpath exchanges stimulate innovation, enable fast and flexible lightpath set-up, and impose no restrictions on the users or the content.

The open exchange model is particularly well suited for the research world, where there is a strong need for flexible and affordable connectivity. The open exchange model also benefits link owners, as this model allows them to remain in control of the use of the link, without any restrictions imposed by aggregator networks which they would otherwise have to interconnect with.

As the reach and the flexibility of open exchanges increases, more and more research will be able to benefit from this model. The emergence of dynamic multi-domain lightpaths, enabled through dynamic open lightpath exchanges, will create further opportunities for research.

Virtualisation will lead to other ICT resources, including computing and storage facilities, becoming more dynamic in the future. The combination of these dynamic ICT resources with dynamic lightpaths will create the next step towards a generic e-Infrastructure for research.

Annex I GLIF

"The Global Lambda Integrated Facility (GLIF) is an international virtual organisation of NRENs, consortia and institutions that promotes lambda networking. GLIF provides lambdas internationally as an integrated facility to support data-intensive scientific research, and supports middleware development for lambda networking. It brings together some of the world's premier networking engineers to develop an international infrastructure by identifying equipment, connection requirements, and necessary engineering functions and services. More information is available on the GLIF website at http://www.glif.is/ "