Integrated and Ubiquitous e-Infrastructures with Open Exchanges

Authors (in alphabetical order): Erik-Jan Bos[1], Lars Fischer[1], & Josva Kleist[1]

This paper presents the important concept of open exchanges, and makes the case that open exchanges are a crucial instrument for facilitating global R&E network connectivity. It further makes the case that open exchanges are an ideal vehicle for integrating a wide range of platforms and digital resources into an integrated and ubiquitous e-Infrastructure.

We argue that the establishment of an open exchange in the Middle East is crucial, as it will allow the networks in the region to connect, to integrate other e-Infrastructures, and to create a platform for other world regions to connect to the Middle East. The result will be that the Middle East increases the potential for further becoming a hub for e-Infrastructures and becoming a better connected region for research and education.

1. Introduction

The idea of exchanges is nearly as old as the Internet itself. Initially facilitating the peering of IP networks, the concept evolved into open exchanges as hubs for global R&E connectivity. Networks such as NORDUnet have used exchanges as a key building block in global networks, and they have been used as nexus for collaborative efforts to connect world regions.

In recent years, open exchanges have evolved beyond networks and have become vehicles for integrating a wide set of e-Infrastructures.

2. Open Exchanges and R&E Networking

In the world of Internet, an "exchange" basically is a location and a piece of equipment that allows two or more independent, autonomous networks to exchange traffic. Internet exchanges in particular allow networks, or networks of networks to announce their *autonomous systems* (AS's) to each other, and *peer* IP traffic. Exchanges that bring together many networks greatly improve the efficiency of the Internet, and have facilitated the growth of *settlement free* peering.

Many R&E networks connect to internet exchanges. However, R&E networks work in a more collaborative manner, and exchange not just IP traffic but also interconnect circuits. Realizing that a similar mechanism to internet exchanges would facilitate this, R&E networks since 2000 have operated what became known as *open exchange points* (OXP)[2]. Initially, OXPs were used mostly to allow R&E networks to stitch together high-capacity circuits for demanding science projects. The GLIF collaboration[3] has been the vehicle of many such projects, and since 2001 has defined the OXP concept and a range of technologies that go with it.

OXPs have been created throughout the world, often at places where substantial infrastructure comes together. NetherLight in Amsterdam, GÉANT Open in London, MANLAN in New York City, and many more are well placed to facilitate connections both from within their region and to other regions. In this way, OXPs have been able to attract at great number of connectors, and are today hubs for research and education networking, and for connections between world regions.

OXPs are simple – they are switchboards, allowing those present at the exchange location to connect to each other if they wish. In this way, an OXP such as NetherLight (operated by the Dutch NREN SURFnet)[4] is an important resource for European R&E networks, i.e. a shared resource they use to connect to each other and to networks from other world regions that connect there. Furthermore, networks throughout the world can connect to all of Europe by using this resource.

---

[1] NORDUnet A/S, Kastrup, Denmark.
[2] Freek Dijkstra and Cees de Laat: *Optical Exchanges*. University of Amsterdam, 2004. Available at http://delaat.net/pubs/2004-c-3.pdf.
[3] http://www.glif.is/
[4] http://www.netherlight.net/

In this way, NetherLight is an important instrument for connectivity between Europe and the rest of the world – a building block for world connectivity. By creating and operating NetherLight, SURFnet – on a cost recovery basis – offers R&E networks throughout the world a means for better connectivity that all can share and use.

The major reasons for the success of OXPs, and their ability to facilitate interconnection of R&E networks, is their open nature, and the fact that they are a simple, shared resource. Open means that anyone can connect to the OXP. This allow any number of networks to come together in a single place, and each make their own decision about who connects to whom, without the need for transit through other networks, or the need for complex policy. It is important to realize that this does not imply that networks give up their autonomy. On the contrary, connecting to an OXP gives the connector the *possibility* to connect to those present at the exchange, not the obligation to do so. Connecting to an OXP does not mean that one must connect with all other networks there; a connection between two parties at an OXP happens only when both explicitly consent.

3. The NORDUnet Global Network and Open Exchange Points

NORDUnet is a regional R&E network that serves the five Nordic countries in Europe: Denmark, Finland, Iceland, Norway, and Sweden. Each country has a national R&E network (NREN), serving scientists, educators, students, and their institutions in the country. Since the 1980s, these five NRENs have worked together to connect their network infrastructures to each other, and to jointly connect to the rest of the world. The NRENs formed NORDUnet as a company to provide regional and global connectivity. Working together as a region, has allowed the Nordic countries to have first-class inter-connections, and to create global connectivity, that each of the Nordic NRENs would not have been able to build alone.

Today NORDUnet offers a state-of-the-art optical and IP network in the Nordic region. Also, NORDUnet helps the Nordic NRENs support a range of services, often jointly produced for efficiency. To allow the Nordic R&E community to work with global partners, the NORDUnet network expands to reach key sites in Europe and well as in North America. Global connectivity is key to facilitate global science, global research, and global education.

During the past 10 years, NORDUnet has focused its network strategy on open exchange points. All connectivity to regional networks, in Europe or globally, is created through major OXPs. OXPs are used as hubs in the network, and OXPs are the places where we expect to connect with other R&E networks, with the networks of a region, and increasingly also with industrial partners and cloud providers.

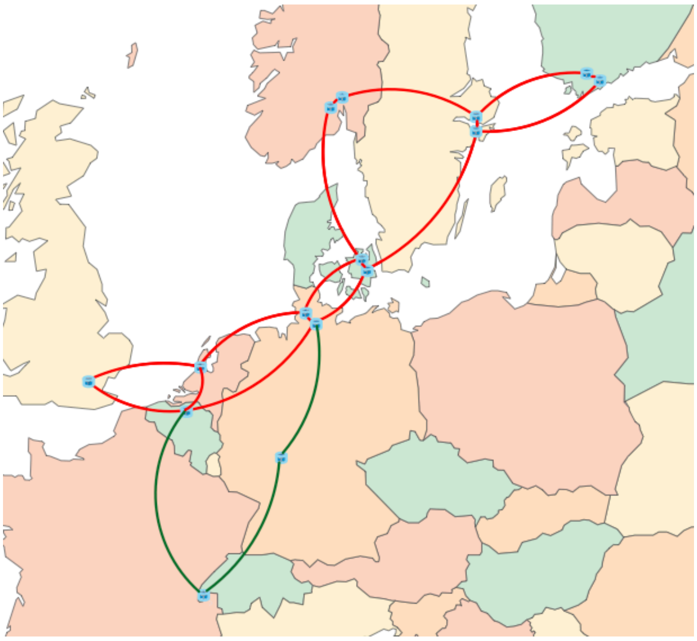

Figure 1. NORDUnet European network.

In Europe, NORDUnet has built a backbone (see figure 1) that, throughout the Nordic region and northern Europe, connects to major European OXPs such as NetherLight in Amsterdam and GÉANT Open London and to Internet Exchange Points such as LINX and AMS-IX. Combined with the establishment of OXPs at the edge of the NORDUnet network, such as NOX-HEL in Helsinki[5], this allows NORDUnet to build a network across northern Europe that connect to partners in the east, west, and south through OXPs.

A recent expansion of NORDUnet's European network to reach CERNLight was created in a similar way. As other European networks, notably SURFnet and PSNC/PIONIER, have adopted a similar strategy, NORDUnet has been able to partner with these networks to jointly connect to OXPs in northern and central Europe, increasing connectivity and adding resilience while reducing cost.

In 2009, NORDUnet decided to establish presence in North America, by extending its backbone and creating North American points of presence. Traffic with North America continued to grow, e.g. as a result of increased global use of advanced instruments in a shared way, and the Nordic research and education community were increasingly using private cloud services such as storage and learning management services. It was critical for NORDUnet be able to offer guaranteed connectivity. When expanding to North America, NORDUnet connected to OXPs such as WIX in the Washington D.C. area, and MAN LAN in New York City, creating trans-Atlantic connections from European OXPs to OXPs in North America. In this way, NORDUnet could directly connect to all R&E resources present at each of the OXPs. The strategy allowed NORDUnet to further expand, connecting from OXPs on the North American east coast to reach OXPs in other regions of the USA (see figure 2), connecting not only to North American based networks, but also to networks from South America, Asia, and the Pacific that connect into these OXPs.

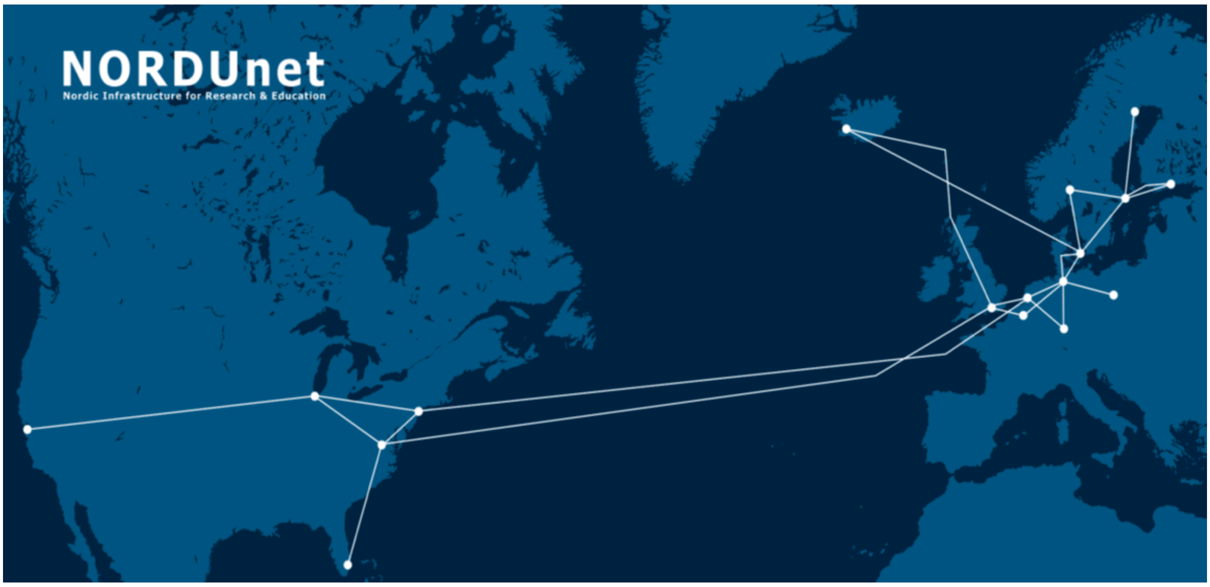

Figure 3. NORDUnet global network.

For NORDUnet, OXPs are a key resource for connecting to the world. We create OXPs at the edge of our network to allow other regions such as the Baltic States and Russia to interconnect, and we build networks to reach OXPs in other regions, to connect to all networks there. For NORDUnet, an OXP is the demarcation point between regions and between networks. The use of OXPs allows us to connect to all present networks and resources that desire to peer with NORDUnet, not just one single network. Direct peering avoids issues of transit and the accompanied (sometimes severe) policy issues, as the direct peers can themselves determine which routes to exchange, instead of relying on and being at the mercy of the transit network. It also allows connections to private, industrial resources in a region. This makes connectivity to an OXP very cost effective.

4. The Advanced North-Atlantic Network Collaboration

In 2013, NORDUnet was faced with a serious challenge from continued traffic growth, also on links across the North Atlantic ocean. NORDUnet had several trans-Atlantic 10Gbps links, and the cost of continuing to add more was prohibitive. We knew we had to move to 100Gbps, but that posed three significant problems: 1) We had to convince the market to offer full channel 100Gbps trans-Atlantic links, 2) we had to persuade them to offer such links at a reasonable price, and 3) we would ideally need three 100Gbps links for resilience. And still, the cost would be prohibitive.

[5] https://www.nordu.net/article/nordunet-expands-its-open-exchange-point-presence-introducing-nox-hel

This was an opportunity for partnership. The Advanced North-Atlantic (ANA) collaboration was formed, consisting of SURFnet, Internet2, CANARIE, ESnet, GÉANT, and NORDUnet. Together, the partners managed to convince subsea cable operators to offer 100Gbps at a reasonable price, and until today to acquire three separate 100 Gbps links, on geographically diverse routes (see figure 3). This system is now called ANA-300G, allowing the partners to share expensive links while remaining autonomous.

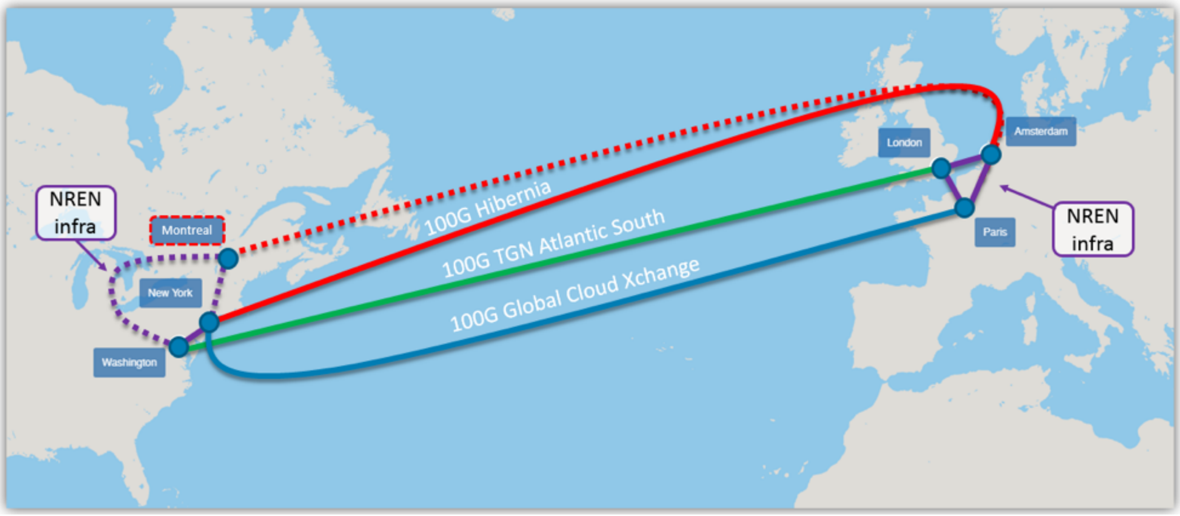

Figure 3. ANA-300G (also showing the imminent move of the northern link to the Montreal OXP.

These separate links now had to be connected on both sides of the North Atlantic. The landing points had to be interconnected with terrestrial links to form a full system. And the system had to be connected to the collaborators' networks on both sides of the North Atlantic, in a way that allowed all partners to connect and use the full system, and have the benefits of full resilience.

To meet these requirements, OXPs were again used. OXPs are key components of the ANA-300G system. All ANA-300G's trans-Atlantic links terminate on OXPs, and all OXPs on a continent are interconnected with links provided by the partners, sometimes dedicated links, sometimes part of an existing link with enough headroom. All ANA collaborators connect to ANA-300G at OXPs. The open nature of an OXP allows all partners to access the system. OXPs are instrumental for the collaboration; the collaboration happens at the OXPs and this is where partners use the shared resource.

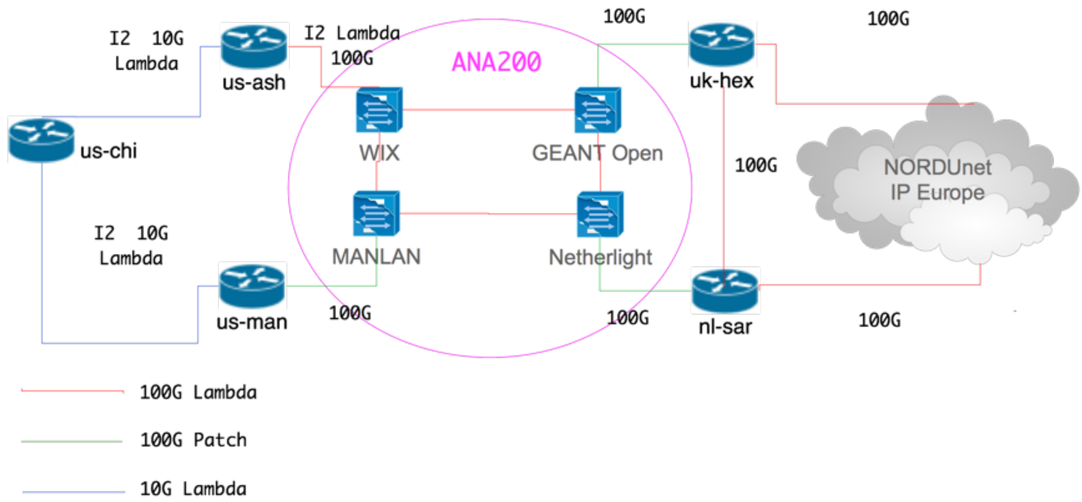

Figure 4. NORDUnet's North-America and European networks interfacing to the ANA system.

OXPs are also the basis of operations of ANA-300G. As links are directly connected to OXPs, to avoid adding expensive equipment in the path, OXP operators also manage the links and facilitate operational procedures. The entire system is operated in a federated manner, with the operators of the OXPs collaborating as a joint operations unit for the system.

In this way, OXPs have allowed the ANA-300G collaborators to build a cost-effective shared system for both trans-Atlantic capacity and continental network access, sharing network resources as well as operations. For NORDUnet, this approach is an extension to the OXP-based strategy already in use. NORDUnet can interface to the system at OXPs in Europe and North America, in fact at OXPs where we were already present for other reasons (see figure 4).

5. The Global Network Architecture

Research and education are today global endeavours. Students and educators work on global subjects. Researchers work in global projects, forming global teams. Research depends on global resources such as unique instruments and facilities. As a result, R&E networks are faced with the challenge of facilitating global access, global connections, and global collaboration. For this rason, the network is sometimes seen as part of the scientific instrument.

Creating global networking is too big a task for one or just a few entities to take on, and too complicated for a single collaboration. Also, if each region tries to connect to any other region of the world, this leads to suboptimal and costly solutions. However, for access to be truly global and truly ubiquitous, in an efficient and cost effective way, a level of coordination must be in place. There is a need for guidelines on how to interconnect regions.

The Global Network Architecture[6] (GNA) is an endeavour to reach consensus, on a global scale, for an ambitious and all-inclusive blueprint for a global network interconnect for research and education. The GNA envisions a global environment, where each region is networked, and the regions connected to each other with inter-region (often trans-oceanic) links. Open exchanges are at the heart of this vision (see figure 5).

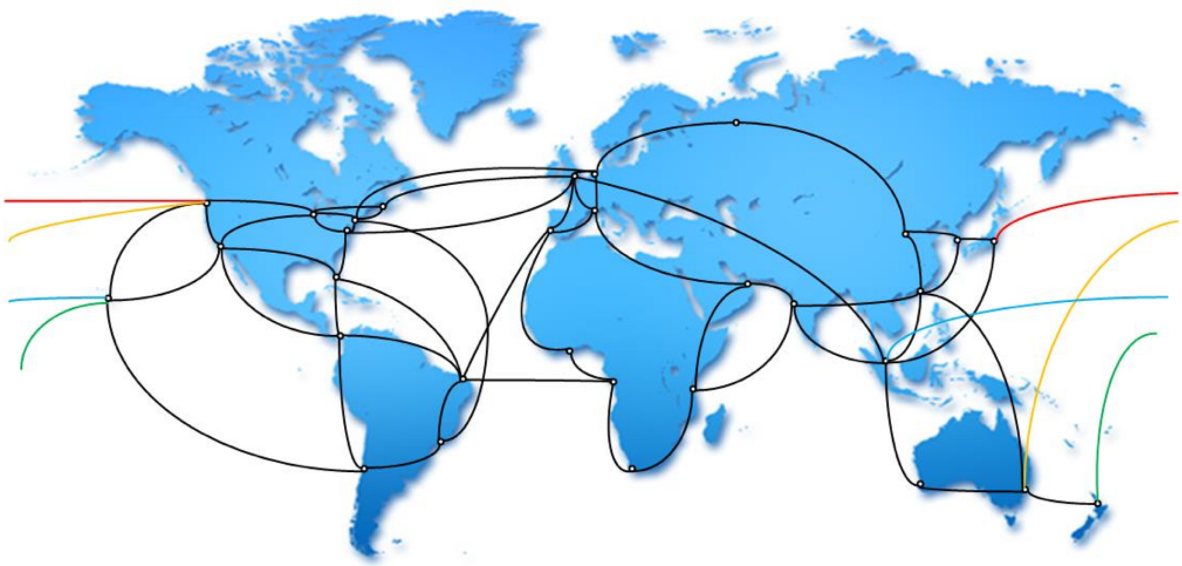

Figure 5. Artist's Impression of a global network interconnect for R&E along the thinking of the GNA.

In the GNA, each region has a few key OXPs that are focal points for R&E networking in the region. National and regional networks connect to these OXPs, thereby connecting all R&E networks, users, and institutions in the region to the OXPs and to each other. In this way, OXPs are the connection points of entire geographical regions and continents. Between the regions, GNA inter-regional systems (such as ANA-300G[7]) are projected to exist. As interregional projects adhere to the commonly agreed blueprint of the GNA, interworking, e.g. on services, is a given.

OXPs play a crucial role in the GNA thinking. OXPs are where a region connects. It is where all the networks in a region can interconnect, both directly and through a regional network. It is also where you can connect *into* a region. By creating a link from another region to an OXP, one connects not just to one site or one network, but to all the R&E networks and to all the users in that region. By creating inter-region systems connecting to multiple OXPs at both ends, full diversity and resilience can be realized for all participants,

[6] http://gna-re.net/resources-downloads/

[7] In fact, the ANA-300G system is considered a pathfinder for GNA. Other pathfinders exist, such as the creation of OXPs in Canada and South Africa, and the addition of links to other world regions from the OXPs that participate in ANA-300G.

allowing full inter-regional connectivity, while maintaining full autonomy (as demonstrated by ANA-300G). In the GNA vision, OXPs are gateways to a region, connecting users and connecting users to resources across the globe.

In the GNA vision, OXPs are key building blocks for a connected, cost effective, and resilient R&E networking world. OXPs are also key contributions to a connected world. By offering an open exchange, a network or an institution can help promote connectivity within a region and enable connectivity to that region from other regions. Using a few, well-placed open exchanges, a region can create both the basis for integrated e-infrastructures in the region, and a platform for making the region part of the GNA in an efficient way.

6. Global Peering Collaboration

NORDUnet, like many R&E networks, connects to Internet Exchange Points like LINX, AMS-IX, and PAIX in California to peer with commercial ISPs. By pursuing an active peering strategy, NORDUnet is able to satisfy over 90% of its traffic with the commercial Internet and service providers through settlement-free peering. Hence, NORDUnet needs to buy transit for less that 10% of its commercial traffic – and is even able to procure such transit at favourable prices by being present at a number of Internet Exchange Points in Europe and in North America.

As other R&E Networks pursue an active peering strategy, NREN collaboration to maximize benefit is natural. The plan for a Global Peering Collaboration is to actively coordinate peering routes among R&E networks on a worldwide scale. This will allow participants in the collaboration to further drive the settlement-free peering towards 100% of the off-net traffic, reducing cost.

Currently, the Global Peering Collaboration is a pilot in which NORDUnet collaborates with Internet2 in the USA and AARNet in Australia, in a proof-of-concept set-up. OXPs are at the heart of the collaboration, as exchange of peering routes between the partners require facilities for exchanging traffic between all the partners, and major links between these facilities.

7. Delivering Global Computing and Storage Resources with Open Exchanges

The importance of OXPs goes beyond network connectivity with and between regions. Global research and education have excellent and ubiquitous network connectivity as prerequisite. Researchers work in global collaborations and form global projects. These collaborations produce data that must be shared and exploited globally. Global science is often also open science, allowing open access and the use and re-use of results, in order to enable novel research to use data in new ways.

Likewise, research facilities share resources globally. Some instruments are so expensive that they can only be afforded through a global partnership. Examples include the Square Kilometre Array[8] (SKA) and the Large Hadron Collider[9] (LHC). Facilities for using the data from instruments, such as high performance computing resources, are also costly and must be shared to be justified. Databases with research data in a certain area can be enormous and are unique. Global access to the research data is often an important requirement in these projects.

With a networking model based on OXPs, we can facilitate access and sharing by connecting key resources to the participating OXPs in the region where they are located – either directly or through an NREN or regional network. With architectures like the GNA, with regional networks such as NORDUnet and systems such as ANA-300G creating resilient connectivity between OXPs, global access to OXP-connected resources is assured.

As an example, a major instrument like the LHC at CERN is also home to an OXP in Geneva, called CERNLight, allowing regional and global networks to connect to each other and to CERN. In the LHCONE[10] overlay network, used to connect physicists at sites globally to CERN and to each other to share LHC data, OXPs are used to connect the various regional sub-networks that together make up the LHCONE network. A similar model can be used for facilitating access to major HPC centres used for scientific computation, or for interconnecting major storage facilities for research. As an example, the A*Star storage facility in Singapore

[8] https://www.skatelescope.org/project/
[9] http://home.cern/topics/large-hadron-collider
[10] http://lhcone.web.cern.ch/

has demonstrated[11] interconnections with facilities and users in North America and Europe by connecting A*Star to an OXP and using inter-OXP links to extend the Infiniband storage protocol globally.

In this model, NRENs have two essential roles:

1. NRENs connect resources such as instruments, databases, and HPC facilities, to one or more regional OXPs, enabling access to these resources, and
2. NRENs connect all their researchers and students, and all their institutions, to an OXP, allowing them to access OXP-connected resources globally.

For this to work, it is critical that users and resources are connected to OXPs, and it is critical that OXPs are inter-connected. The inter-connection of OXPs can be done by regional R&E networks such as NORDUnet and GÉANT, or sometimes as part of inter-regional consortia such as ANA-300G, where partners provide both the inter-regional links and the links connecting OXPs in the region.

The use of OXPs allows R&E network services to be integrated and to become ubiquitous in a truly global way. It allows a wide range of infrastructures to be connected in an efficient manner, and allows open access that facilitates integration of resources from many institutions and organizations, forming a globally federated e-infrastructure platform. OXPs allow this to be done with a minimum of administration or policy overhead. OXPs form a natural perimeter, and serve to set a useful and well-understood division of responsibilities.

When users, resources, and regions are interconnected at open exchanges, connections are created using an open, transparent facility, with full autonomy for the connector. Each connecting party controls their own connection, and therefore decides how and with whom to interconnect. Any policy is the sole decision of the connecting institution or network. Operations, cost, and uptime are also the responsibility of the connector. Each connecting network and each inter-regional connecting OXP have one job.

With OXPs, we can have ubiquitous, global access for anyone, anywhere, anytime, without the need for a complicated regional or global governance, and without the added complexity of transit rules and transit policies.

8. Integrating Public and Private Cloud Resources with Open Exchanges

Recent years have seen a dramatic growth of the use of cloud storage and cloud computing in research and education – and indeed throughout society. Use of clouds for research and education has been multifaceted, arising from many use-cases and driven both top-down and bottom-up. Some use arises from individual researchers deciding to use cloud resources, or a small research group deciding to use cloud facilities for sharing research data. Other uses are created by large institutions or government initiatives, sometimes with a focus on the financial efficiency, sometimes with a focus on cloud capabilities for sharing and collaboration.

In a similar way, e-infrastructures have offered cloud facilities in a multitude of ways. This ranges from simple desktop storage synchronization services[12] bought from commercial vendors by individual researchers or procured collectively by institutions or NRENs, to private cloud storage and computation operated and offered at institution or national level, to major use of commercial cloud infrastructure-as-a-service offerings.

The result is a great diversity of services, products, applications, and use cases, and an important requirement on NRENs to enable them all. As research and education are by nature global, science collaboration and open science require us to facilitate access to these cloud resources for researchers and students, no matter where they are. Unlike many business applications, we cannot allow an assumption of local use.

In Europe, the European Commission has launched an initiative called the European Open Science Cloud[13]. The objective is to create an open, comprehensive, integrated, and ubiquitous e-infrastructure spanning all of Europe, allowing all European research data to be discovered, shared, combined, and used by researchers anywhere in Europe. The vision is an integration of institutional, national, and pan-European cloud resources, applications, and computational facilities.

---

[11] http://www.geant.org/News_and_Events/CONNECT/Pages/infinicortex_project.aspx
[12] I.e., Dropbox, Box.com, etc.
[13] http://ec.europa.eu/research/openscience/index.cfm?pg=open-science-cloud

CERN has launched initiatives[14] such as Helix Nebula to allow storage and computation for the LHC to make use of cloud resources, for flexibility, dynamic growth, and cost efficiency. These initiatives are focused on procurement of cloud storage and cloud computing engines from commercial vendors, and on integrating these into a scientific computing environment.

To integrate these diverse resources, of any scale and any type of ownership, and to allow them to be openly accessed by users anywhere, at any time, we need a simple and solid architecture, allowing simple cost sharing and reducing the need for (often complex) policy and governance. We argue[15] that best practice for doing this in an efficient and cost effective way is to use open exchanges. Due to their open nature for connecting, open exchanges can connect any kind of resource, commercial or private, either directly or through an intermediary network. And with regional and inter-regional networks connecting to an OXP, once a resource is connected to an OXP, it is connected to the world. Using OXPs in a model much reminiscent of the GNA, we break down the policy, governance, and cost issues of providing open access to cloud resources (see figure 6), giving us a straightforward strategy for connecting cloud resources of all kinds and sizes.

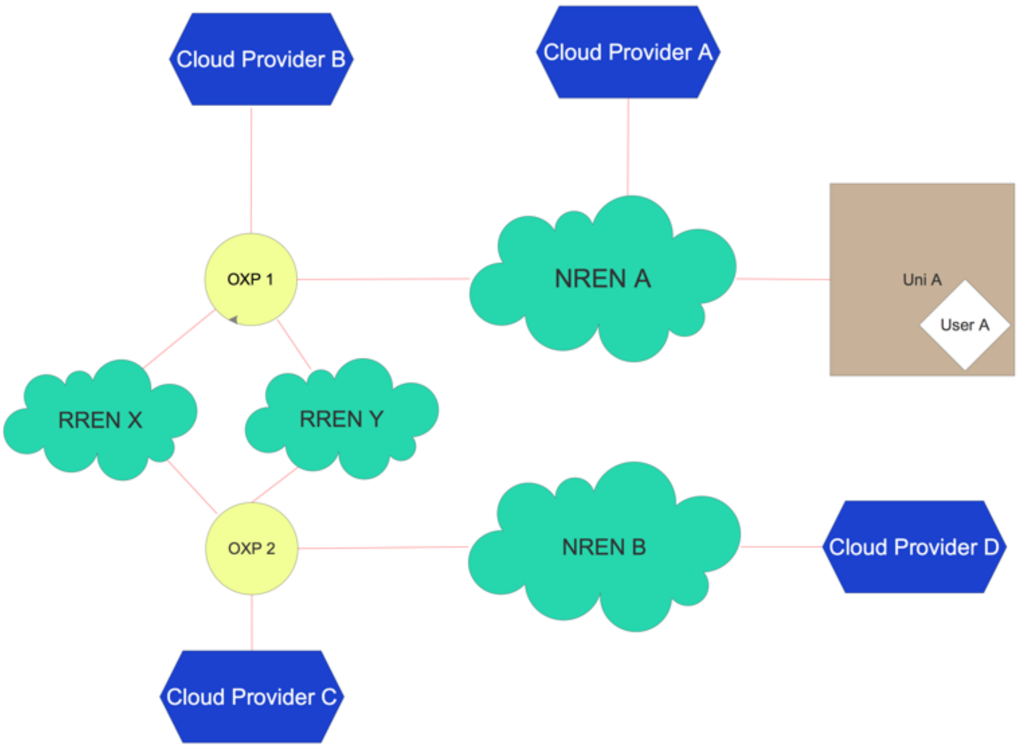

Figure 6. The generic model for connecting to cloud providers for R&E networks.

In this way, OXPs become the focal point for sharing of cloud resources, and for sharing data and capabilities in the cloud by research and education. OXPs are the key instrument for integrating cloud resources into a global, ubiquitous e-infrastructure. OXPs also become the focal point for connecting the R&E community to commercial providers of cloud resources. In fact, most major cloud providers globally have facilities near one or more OXP, allowing them to be directly connected to an OXP.

9. The Case for OXPs in the Middle East

The open exchange concept originated in Europe and North America, but in recent years has spread to other world regions. OXPs have been established in Asia (Singapore, Hong Kong, Seoul, Tokyo), South America (Santiago de Chile, Rio de Janeiro), and Africa (Cape Town). This demonstrates the acceptance of the idea of building global e-infrastructures based on OXPs. With more OXPs being created throughout the world, more and more world regions can be connected by building infrastructure between OXPs.

[14] http://www.helix-nebula.eu/
[15] https://www.nordu.net/content/best-practices-cloud-provider-connectivity-re-users

For the Global Network Architecture (GNA) to be truly global, it is important that all world regions have open exchanges, so that they can be connected to the global platform being created. Likewise, we believe OXPs are essential to creating an integrated ecosystem of diverse e-infrastructures, as argued above.

For these reasons, we believe it would be of great benefit if one or two OXPs were created in the Middle East region. Several benefits would follow:

- It would allow networks throughout the region to connect to a few, shared resources, creating the possibility for flexible interconnects.
- It would allow projects in the region to use such OXPs as basis for creating shared connectivity to other world regions, following the model of the ANA-300G collaboration – possibly in conjunction with other world regions.
- It would allow projects building advanced connectivity between other world regions to connect also to the Middle East. A great number of cable systems between Asia, Europe, and East Africa pass through the region. Projects building connectivity between the Asia Pacific region and Europe could *break out* a connection to a well-positioned Middle East OXP at little or no extra cost. As a result, such an OXP would be well connected to China, Asia, the Pacific region, Africa, and Europe, greatly improving connectivity options for all networks in the region.
- It would improve reach and efficiency of global R&E networking. If projects connecting Africa, Asia, the Pacific, and Europe come together in a Middle East OXP, they will offer the option of connecting all the regions in an any-to-any manner, rather than just connecting between a pair of regions, greatly improving value.
- It would allow integration of e-infrastructure platforms in the region and globally into a ubiquitous e-infrastructure for the region.

For these reasons, we would much welcome local initiatives to establish and share such an OXP with a fair and affordable connection policy. It would be a highly valuable resource for the global R&E community, and we believe it would benefit all, both locally and globally.

10. Conclusions

In this paper, we have outlined the history of open exchanges (OXPs), shown how the concept has played a crucial role in the expansion of the NORDUnet network, and through the examples of global peering strategy, the Global Network Architecture (GNA), and the ANA-300G collaboration, shown how it is becoming essential to both regional and global R&E networking and network collaboration.

We have argued that OXPs are furthermore an essential building block in forming a globally connected, integrated, and ubiquitous e-infrastructure for research and education. OXPs play a central role in connecting and integrating a diverse set of resources from public and private actors, and delivering such infrastructure to researchers, educators, and students everywhere.

Finally, we believe that there would be great benefit – both for the region and for global R&E networking as a whole – in establishing one or more open exchange points in the Middle East.