

Best Practices¹ for Cloud Provider Connectivity for R&E Users

Authors (in alphabetical order): Erik-Jan Bos², Lars Fischer², David Foster³ & Josva Kleist²

Date: 31 August 2016

Version: 2.0

1. Introduction

R&E Networks have been in the business of serving the needs of research and education for decades. A recent development is that more and more R&E Networks are required to support the overall business of research and education for their customers. As R&E institutes have felt the pressure from governments to become more efficient and cost-effective, their interest has turned to cloud solutions for scientific applications as well as back-end office systems.

The use of clouds, both commercial and private, is increasing rapidly. Large scale connectivity with cloud providers is a rather new but important area, in which R&E Networks are trying to find their way to add value⁴. Connectivity with commercial cloud providers nowadays is an important topic, and it is becoming crucial that advice to policy makers, decision makers and procurers⁵ is given so that over time it will lead to a coherent, scalable and increasingly cost-effective solution for connecting to cloud service providers.

2. Purpose of this paper and audience

This paper is a "Best Practices paper" that seeks wide adoption and agreement in the R&E community. It is intended as a reference for policy makers, decision makers and procurers⁴ of cloud services giving guidance in terms of the network architecture that should be adopted by cloud providers.

It is important to provide clear advice on what should be specified as network connectivity requirements in the general case. This will help promote the coherent use of the R&E networking infrastructure that is in place. Furthermore, it is consistent with the goals of the European Union's Digital Single Market and similar initiatives around the world for service provision.

The net result of adopting these best practices should be a reduction in uncertainty around network solutions and, over time, reduced cost in producing responses to calls for tender.

The approach proposed does not preclude requests for direct connectivity, typically a high-speed point-to-point connection, between a cloud provider and a user site in special cases.

¹ This paper is written on an *ad personam* basis, and the authors are thankful to many individuals from all over the world that spent time and energy on commenting earlier versions of this document. Their input helped tremendously to get this paper into its current shape.

² NORDUnet A/S, Kastrup, Denmark.

³ CERN, Geneva, Switzerland.

⁴ A recent proposal to create a Cloud Provider VRF is deemed not-scalable by the authors of this document as you cannot create a VRF for each purpose. Moreover, seeing each prefix many times at many peering points does not scale for large R&E Networks.

⁵ Persons that write an RfP and run the procurement.

3. Related Work

As cloud connectivity is a current and key issue for scientific computing, a number of stakeholders are active in this area. In this paper we try to present what we believe is the current best practice, leading to simple solutions with straightforward governance.

An overview of proposals and work in this area can be found in the Cass & Martelli paper⁶. Note that the *Best Practices* described here has some similarities with the *Cloud Exchange Point* approach described by Cass & Martelli, but is simpler.

4. Terminology

Through-out this paper, terms are used that have a specific meaning in the context of this paper. For the avoidance of doubt, they are listed below in alphabetical order:

- AUP: Acceptable Use Policy. An NREN typically works under an AUP, prescribing what type of traffic is allowed on the network, and what is not. A regional network, such as GÉANT or NORDUnet, and has no AUP by itself but typically works under the collective AUP of its member NRENs.
- Cloud: Computer and/or storage infrastructure spanning multiple servers and often multiple physical locations.
- Cloud Provider: A commercial or private organization delivery Cloud Services. Commercial examples include Google, Amazon, Microsoft, and Box. Private examples include the EGI Federated Cloud, and GRnet's Okeanos.
- Cloud Services: The collection of storage and compute services running on another organizations infrastructures and accessed through the Internet.
- EOSC: European Open Science Cloud⁷.
- GNA: Global Network Architecture, the coordinated effort of leading R&E Networks to define a blueprint for a next generation global interconnect for research and education⁸.
- LHCONE: Large Hadron Collider Open Network Environment⁹.
- NREN: National Research and Education Network. A network service provider serving institutions for research and education in a country. Services are typically not limited to network services only. Examples of an NREN include GARR, SUNET, and Belnet.
- OXP or Open Exchange Point: A carrier neutral data exchange facility, connecting R&E Networks, commercial ISPs, and Cloud Providers. Examples include NetherLight, CIXP, but also commercial exchanges such as AMS-IX and LINX.
- R&E Networks: The name used to denote the collection of NRENs and RRENs (see below).
- RREN: Regional Research and Education Network. A network service provider for R&E serving a region typically larger than one country. A RREN is a provider for NRENs. Services are typically not limited to network services only. Examples include GÉANT and NORDUnet.

⁶ Tony Cass & Edoardo Martelli: Data-Intensive Cloud Service Provision for Academic Institutes: the Network Connectivity Problem. CERN, August 2016.

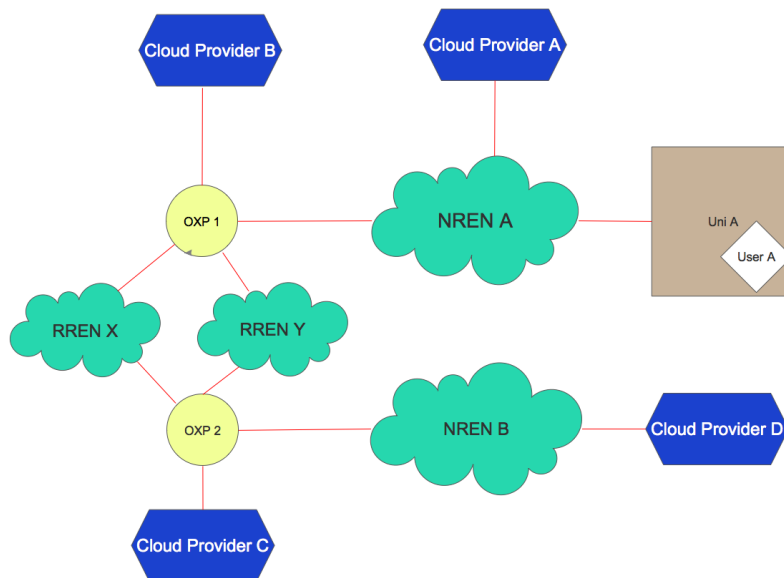
⁷ More information on EOSC can be found at: <https://ec.europa.eu/research/openscience/index.cfm?pg=open-science-cloud>.

⁸ More information on GNA can be found at <http://gna-re.net/>.

⁹ More information on LHCONE can be found at: <http://lhcone.web.cern.ch/>.

5. Reference Model

The model used throughout this paper can be summarized in one picture, which will be discussed in the following text:



6. Issues and Breaking Them Down

It is possible to break down the issue of connecting R&E users to private and commercial cloud resources, and ensure seamless interconnectivity of such resources, into a few easily manageable tasks. Open exchanges are key to achieving this in a flexible, cost-effective and transparent manner.

The approach is that all entities – research and education organizations with their users, user resources, and commercial providers of storage and compute resources – are interconnected directly or indirectly through open exchanges and that open exchanges are used for providing both resource-to-resource and user-to-resource connectivity.

In the simplest example, a campus-based user is making use of resources at Cloud Provider B. If both Cloud Provider B and the NREN serving the campus are connected to an open exchange, the NREN can facilitate the connectivity needed. In a slightly more complex example, resources from two providers are used. In this case, if all three parties are connected to the OXP, traffic on behalf of the user both between the campus and the providers as well as between the providers can be facilitated over the OXP.

When more than one OXP is used, network trunks between the open exchanges must be provided. This is similar to the approach adopted by the GLIF¹⁰ Community and in the GNA.

The advantage of this approach is that rather than looking at the task of a fully meshed interconnect of user institutions and providers, and the cost-sharing and policies of such a mesh, we can address three separate tasks:

1. Connecting a user institution, typically through an NREN, to an open exchange,
2. Connecting a cloud provider (commercial or private) to an open exchange,
3. Creating trunks between open exchanges.

¹⁰ More information on GLIF can be found at <http://www.glif.is/>.

Using these three components, all required connectivity and use cases can be served. Some readers will note that this model is close to the original design proposed for LHCONE, based on open exchanges.

One extra advantage of the model proposed in this paper is that it is a general model for the R&E institution – Cloud Provider connectivity. It will work equally well for e.g. cloud services which may be part of EOSC.

7. Requirements for this to work

If we address the tasks outlined above, we note that:

1. *Connecting a user institution to an open exchange* is in most cases trivial. Most institutions are already connected to an NREN that connects to a OXP, either directly or through an RREN. In a few cases (such as CERN), institutions also connect directly to an OXP, next to their NREN connection(s). In either case, network capacity and required services are in place, and are largely covered by existing costing and cost-sharing agreements.
2. *Connecting Cloud Providers to Open Exchange Points* has started during recent years. In fact, for some providers such connections are already in place, e.g. Amazon at GÉANT Open London, and Microsoft and others at NetherLight. For cloud providers not already connected, options exist to easily make this happen and the cost of this can be directly assigned to the cloud provider's services (or not, and seen as sunk costs), i.e.:
 - a. The cloud provider comes to the OXP: The cloud provider has an internal, proprietary network and will already be in locations served by one or more OXPs. Connecting this provider directly to the OXP is a fairly trivial task.
 - b. The cloud provider uses the NREN to come to the OXP via Layer 2: The cloud provider identifies an R&E Network that can provide connectivity between one of his PoPs and an OXP, as a layer 2 or 2.5 network service. Or, the cloud provider can consider the use of a local or regional telco to supply this connectivity.
 - c. The cloud provider uses the NREN to come to the OXP via Layer 3: The cloud provider may possibly connect to the IP service of an R&E Network and this will allow the provider's traffic with the OXP to flow over shared IP, possibly using a tunnel or a VPN.
3. *Inter-exchange bandwidth* will have to be acquired and needs to have sufficient capacity for the traffic and distribution of user institutions and providers. This is similar to the case of both LHCOPN and the trans-Atlantic links for LHCONE. Such links can be both mission specific, or general purpose, for all users. In either case, it should be noted that major R&E Networks have already invested in high bandwidth interconnects between locations of OXPs, and are often already using these exchanges themselves to terminate links. Hence, OXP-OXP links for cloud traffic can easily be provided at layer 2 or 2.5 by these R&E Networks. It should be noted that multiple links, from different sources, can exist between any pair of OXPs, adding to bandwidth and resiliency.
4. Ensure the correct policy is in place. The community and governments need to recognize cloud provisioning is an end-to-end service, and for this to work we need AUPs that allow this for each NREN in the path.

There are a few additional requirements:

- Links and networks used must allow for Cloud <-> Researcher traffic, e.g. for the scientist to store his data at a cloud provider.
- OXPs and any links used to connect commercial cloud providers must also allow for Cloud <-> Cloud on behalf of a researcher, e.g. to ensure data stored at Google can be transferred for processing to Amazon. This is in line with the trends towards data portability between providers.
- Peering with a cloud provider by the R&E Networks and offering these routes to other R&E Networks is important for this to work.

8. Business Case Considerations

The use of OXPs as demarcation for an R&E user-to-cloud provider network architecture serves to simplify the network design into manageable pieces. Funding for the architecture is best approached in the same way. The basic components of the architecture are: OXPs, provider-to-OXP links, OXP-to-user networks, and inter-OXP links. We address the funding of these below.

- OXPs are self-supporting and in most cases already using a cost-recovery model in which the connectors contribute to the cost of operating the OXP. This model is globally accepted and used.
- When the R&E community enters a procurement of Cloud Services based on the notion that traffic is exchanged at an OXP, connecting to an OXP becomes part of the service delivery, and hence part of the procurement. In practice, since OXPs are often located at or near major infrastructure hubs, cloud providers will often already have infrastructure at or near OXPs.
- Connecting R&E users is the business of R&E networks. Most R&E networks are already connected to OXPs, either directly or through an RREN. Each R&E network is funded through a different mechanism, but connecting to critical resources is always part of the basic service model.
- Interconnecting OXPs and serving OXP–OXP traffic is similar to serving traffic between R&E networks, both inter- and intra-continently. This is a service R&E Networks have provided for decades, and have cost-shared, through the establishment of continental backbone and substantial inter-continental links. Serving such traffic flows for their users is the core business for the R&E Networks, and these networks are well positioned to extend this to the model described here, using the cost models and facilities already in place.

Intercontinental transit is by nature an extension of the OXP-OXP area. However, as intercontinental is costly and sometimes limited, more care is needed to ensure fairness and reciprocity, such as is being worked on the ANA-300G Collaboration and within the GNA.

It should be noted that a large cloud provider might choose not to be present at many OXPs on one continent, but instead wants to connect to one or two well-established ones and wants to rely on the R&E Networks to distribute the connectivity. This is fully in line with this paper.

9. OXPs in the R&E Networking world

A cloud provider present at an OXP is assumed to have a connection to the peering fabric. With direct peering over Layer 2 through the OXP peering fabric, the R&E Networks present can peer with the cloud provider directly, or extend the peering to the end user institution, as required. RRENS, such as GÉANT and NORDUnet, present at the OXP can peer with the cloud provider directly, and carry the routes to their constituencies. The use of a simple and effective scheme for BGP Communities will help NRENS decide how to handle the cloud provider prefixes.

In Europe, the major OXPs with and without R&E Network involvement are:

- GÉANT Open London and LINX in London
- NetherLight and AMS-IX in Amsterdam
- CIXP in Geneva
- NOX-HEL in Helsinki
- GÉANT Open Paris and SFINX in Paris

Other major OXPs with R&E Network involvement around the world are:

- MAN LAN in New York City
- StarLight in Chicago
- WIX in Washington, DC
- PacificWave in Seattle, Sunnyvale and Los Angeles
- AMPATH in Miami
- Singapore OXP in Singapore
- Soon: Montreal OXP in Montreal
- Soon: Cape Town OXP in Cape Town

Some countries or regions will not have a local OXP. In this case NRENS and cloud providers can either be connected to the nearest OXP, or work to establish a local OXP can be undertaken.

10. Recommendations and Future Work

We recommend for users, institutions, and projects using cloud resources, including resources from commercial providers, to base their service model on the simple networking model described in this Best Practices document. Specifically, we recommend that the delivery of Cloud Services to R&E users through the NRENs be based on the model described herein.

R&E networks and user communities have been establishing many of the components and network resources needed for implementing this model. In Europe, OXPs have been established by GÉANT, NORDUnet, CERN, SURFnet, and others. Major networks resources that interconnect these have been put in place by GÉANT, NORDUnet, and others. For trans-Atlantic networking, the ANA-300G Collaboration has deployed a system of 3x 100Gbps, interconnecting at OXPs on both sides of the North Atlantic, with a reciprocal back-up agreement with ESnet.

Some OXPs are already connecting cloud service providers, either directly or through an NREN intermediary. For example, SURFnet is connecting the Microsoft Azure cloud to NetherLight, as a subcontractor to GÉANT, and is sharing the connection with European NRENs via GÉANT. We recommend that many more such connections are established. A suitable follow-up would be to analyze this further, and look more closely into the cost and charging models. An outcome could be a document that describes this to the cloud providers, highlighting the benefits of connecting to OXPs to reach the R&E world. In general, this will help to convince them to play ball with the R&E Network community.

Additionally, more work is needed at the network routing and addressing layer. The best practices described here ensure optimal connectivity, but cannot ensure that connectivity is used. To ensure that high-quality R&E networks are indeed used to reach cloud resources, IP-addressing and routing should be controlled. Current work at CERN and elsewhere points to challenges with current (commercial) cloud resource offerings. More work is needed in this area to fully understand the issues at hand, and can later serve as input for future tenders.

This best practices document should be followed by a recommended set of specific requirements¹¹ that can form part of commercial tender documents by providing a model text for procurements for cloud services building on the recommendations of this document. GÉANT has led various efforts on procurement of cloud services in Europe for several years now, and has gained a wealth of experience in this area¹². We believe that GÉANT is well positioned to lead this effort and given the growth of tendering for cloud services we believe this should be produced and widely disseminated with some urgency.

#

¹¹ Such as in areas like VM addressing, IP routing, and security and performance requirements.

¹² More information on GÉANT Cloud Services can be found at: <http://services.geant.net/clouds/Pages/Home.aspx>.