Global Network Architecture (GNA)
WHITE PAPER

Global Virtualization Architecture for the GNA

(version 0.91)

Authors: Lars Fischer[1], Josva Kleist[1], Erik-Jan Bos[1]
Date: 9 September 2016

The Global Network Architecture is evolving into a blueprint for the future, global R&E network of networks. Already pathfinder projects are deploying implementations of the GNA ideas. These pathfinders bring valuable experience, showing that the GNA ideas are sound, and identify further architecture and technology challenges that must be addressed.

This white paper highlights one such challenge.

Initiating Problem: The GNA Inter-Regional Systems

The ambition of the GNA is to create a globally agreed and inclusive blueprint for the construction of a global network interconnect, based on inter-regional network systems created, funded, operated, and used by multiple partners. The method used is to create an aspiration architecture to which all participants can build, using common elements, concepts and nomenclatures. Moreover, the Commons adds a global perspective and global facility to the inter-regional network systems, that brings extra benefits to all participants.

For an inter-regional system, constructed and run by multiple parties, we would like to be able to guarantee each partner certain capacity and network characteristics, according to their requirements and to what they pay for. We would like this to happen across the entire system, not individual links, with full resilience. We would also like to be able to allocate part of the resources to the Commons, again with guaranteed characteristics. We would like to allocate part of the system to specific use cases, demos or tests, on a short-term or long-term basis. Finally, we would like all users to be able to burst while using and sharing unused capacity[2].

We need to be able to account for and report on resource usage, and we must ensure that the total system in visible and transparent for all partners. We must ensure that the total system provides recovery from failure and adds resilience for all partner configurations, and that all configurations have predictable and deterministic network characteristics. Finally, applying changes and provisioning network services must be predicable, also during set-up and tear-down time.

In short, for a successful implementation of the GNA, we need inter-regional network systems that offer each partner a fully resilient system with deterministic behaviour.

Vision: Virtualizing the Inter-Regional Systems

Our vision is to virtualize the inter-regional network system, allowing each participant to control a complete virtual system, with full resilience, and with pre-defined and deterministic characteristics. With such a system, each partner can create their own topology and interconnects, with appropriate characteristics. Furthermore, capacity can be flexibly allocated to partners, with an ability to burst, assuring that partners can share unused capacity in a pre-defined and pre-agreed manner.

---

[1] NORDUnet A/S, Kastrup, Denmark.

[2] In the global network interconnect the sun never sets, hence using bandwidth in a region that is in its night with low usage from the region itself is one of the concepts behind the Commons. You cannot bank unused bandwidth so you better use it smartly if there is low usage from the owner and demand elsewhere.

Furthermore, additional virtual systems could be created for the Commons and for accepted projects. As an example, a virtual system created for an e-VLBI project, would have a guaranteed 4 Gbps capacity *and* a constant line rate and jitter[3], independent of other virtual systems created on the inter-regional network system.

Typical current technologies for approaching these requirements are MPLS meshed networks or fully routed IP networks. While these technologies can implement resilience and present full systems, and can manage some level of bandwidth guarantees, achieving fully independent systems for partners is not possible. Also, end-to-end network characteristics cannot be guaranteed. At the other extreme, Layer 2 technologies such as OTN offer guaranteed characteristics and full independence, but will not allow sharing and will present links, not full systems.

We believe that to fulfil this vision, full network virtualization should and can be provided. The virtualization must be sufficiently powerful to provide full systems, not just links, and to provide different and independent systems to each partner or even user, with guaranteed characteristics. The virtualization must be transparent to the partners and fully instrumented.

Ambition

Our ambition is to create and architecture and supporting technology that allow full virtualization of inter-region systems based on the GNA, allowing management and traffic engineering for all partners, with transparency. This creates a full system view rather than a collection of links, and with guaranteed network characteristics. It is our aim to allow multiple partners to fully exploit a collection of shared resources.

However, our ambition reaches beyond this network application. We believe that a network virtualization and a Global Virtualization Architecture have applications in a wide range of applications in research and education. We are seeing a number of initiatives, such as the proposed European Open Science Cloud (EOSC), that aim to create a comprehensive platform for open science and data driven research. These proposals call for a global infrastructure that allows to do instantiation, is secure, and enables global collaboration. We believe that network virtualization is an important element in such an infrastructure. In the end, our ambition is for a Global Virtualization Architecture to be sufficiently powerful to provide the network substrate for initiatives such at the EOSC, and, fully integrated with the implementation of the GNA, a global, comprehensive, end-to-end infrastructure for data-driven research, offered to research communities on a global scale.

Example: ANA-300G North Atlantic Network System

ANA-300G today is a pathfinder for a GNA inter-regional network system. The ANA Collaboration has proven that such system is efficient and allows partners to contribute and collaborate based on individual requirements. ANA-300G has also identified limitations of the technology currently used. ANA-300G has exactly the challenges outlined above: The desire to provide each partner with a full, resilient, North Atlantic interconnect system with deterministic behaviour.

With full virtualization in place, a participant like NORDUnet can contribute half of a 100G link to the ANA Collaboration, and receive a complete virtual system that is resilient over all the links of the ANA-300G, with guaranteed bandwidth and network characteristics. Moreover, this also allows other partners to use the resources allocated to NORDUnet when not used by NORDUnet. And yet with all partners protected from being impacted by traffic patterns of other partners.

---

[3] Note that a simple VLAN on a 100Gbps link will not do this. The equipment in the path will use internal buffering to provide an average of 4 Gbps, with bursty characteristics. This has proven to overload end systems' buffering capabilities.

Rationale: Why the Virtualization Approach

One may think that attempting to solve the problem of sharing an inter-regional network system through full system virtualization is overly ambitious, or even over-engineering. However, we believe the approach is both possible and necessary.

It is *possible* because most of the architectural ideas and much of the technology needed already exists. A number of projects have defined virtualization architectures that can be applied. GÉANT has established GTS[4], a network virtualization technology and implementation currently used to offer a pan-European testbed generation and support system. We believe many ideas embedded in GTS can be used to defined a Global Virtualization Architecture for the Global Network Architecture, and furthermore that parts of the technology used to create GTS can be used to develop such systems.

Several groups in the R&E networking community have strong experience designing and building such systems. We believe the a "Global Virtualization Architecture version 1.0" can be realized fairly quickly.

We also think it is *necessary* to create such architecture and to build the system. Virtualization of other ICT infrastructure elements used to support global, open science has been happening for several years. Building research e-infrastructures[5] as virtualized systems based on such technology, is already under way, or in some cases well advanced. Comprehensive infrastructure initiatives for open science, such as the European Open Science Cloud, will rely heavily on virtualization. Virtual computer and storage is taken for granted. It is about time we address the network aspects of such regional and global research e-infrastructures.

We have so far struggled to find a comprehensive architecture for creating a trans-ocean system out of a collection of links, and for providing partners in such a system with independent systems and service guarantees. We believe that through the Global Virtualization Architecture, a possible and necessary component of the GNA, we will be able to offer a solution to the immediate problem, as well provide a strong foundation for the future.

# # # # #

[4] The GÉANT Testbeds Service.

[5] This is the term in use in Europe, while it is called cyber-infrastructures in other parts of the world.